

Directed Information Graphs

Christopher Quinn, *Student Member, IEEE*, Negar Kiyavash, *Member, IEEE*,
and Todd P. Coleman, *Senior Member, IEEE*

Abstract—We propose two graphical models to represent a concise description of the causal statistical dependence structure between a group of coupled stochastic processes. The first, minimum generative model graphs, is motivated by generative models. The second, directed information graphs, is motivated by Granger causality. We show that under mild assumptions, the graphs are identical. In fact, these are analogous to Bayesian and Markov networks respectively, in terms of Markov blankets and I-map properties. Furthermore, the underlying variable dependence structure is the unique causal Bayesian network. Lastly, we present a method using minimal-dimension statistics to identify the structure when upper bounds on the in-degrees are known. Simulations show the effectiveness of the approach.

Index Terms—Graphical models, network inference, Granger causality, generative models.

I. INTRODUCTION

RESearch in many disciplines, including biology, economics, social sciences, computer science, and physics, involves studying networks of interacting, stochastic processes. The human brain, stock market, and Internet are some examples. For numerous research problems, it is important to characterize the *structure* of these large networks of interacting processes that elucidates the extent to which the past of some processes affects the future of others. In particular, it can be useful to have a succinct representation of the structure, such as a simple graphical model of the network. For instance, each node could represent a process, with directed edges between the nodes representing directions of influence.

There is a large body of research on developing well-defined graphical representations of networks of random variables. Markov networks, Bayesian networks, dynamic Bayesian networks, and chain graph models are some well-known examples. In these, random variables are represented as nodes in the graph. Edges represent conditional dependence relationships between the variables. These graphical models can be used for arbitrary sets of random variables, but the relationships they show are mutual. These graphs have been successfully used to represent the structure of networks of static stochastic systems. Some applications include object recognition [1],

error-correcting codes [2], cellular networks [3], and medical diagnostics [4].

Markov networks and Bayesian networks in particular represent two different perspectives on the structure of networks of random variables. Markov networks directly represent the dependence between each pair of variables, conditioned on all other variables. Bayesian networks represent factorizations of the joint distribution, so each variable potentially depends on preceding variables, and then the conditional terms are reduced. See [5] for an overview of graphical models.

Networks of random variables do not lend themselves to concise representation of interacting stochastic processes, which are sets of random variables indexed over time. Representing each random variable as a node results in large, cumbersome graphs growing with time. Moreover, such a representation will not aid with visualization of the structure of inter-dynamics of coupled time series. For instance, it could be difficult to see how the past of some processes affects the future of others.

Examples of such networks of dynamic stochastic systems include financial networks, computer networks, social networks, and biological networks. In financial markets, investment strategies are based on predictions of how past activity of some stocks affects the future activity of others. For instance, suppose that the price of corn, a component of chicken feed, strongly influences the price of chicken. Simulated data is shown in Fig 1. Knowing this causal relationship can lead to better investment strategies than simply knowing the two are correlated.

Another example is with computer networks. To investigate a computer attack, for instance, simply knowing the connectivity structure alone is not sufficient to identify the source. For legal proceedings involving warrants and seizures, it is significantly more useful to know the traffic influence structure (see Fig 2). Likewise, many social networks have non-mutual relationship structures. Twitter, email networks within a company, and gossip networks are some examples where the connectivity alone does not convey the structure (see Fig 3).

One last example is from biological networks. The primate visual system functions through a complex feedback network [6]. As light enters the eye, information is sent from the eye to other parts of the brain, where initially low-level processing occurs, such as processing simple shapes. Then information is sent to higher-level processing, such as object and facial recognition. Simply considering correlation between these parts of the brain might suggest information is sent in a feed-forward manner, from the eye to low-level to high-level processing areas. However, it has been shown that visual processing relies extensively on feedback between these areas

This material was presented in part at NetSciCom 2011, ISIT 2011, and CDC 2011. C. Quinn was supported by the Department of Energy Computational Science Graduate Fellowship, which is provided under grant number DE-FG02-97ER25308. This work was supported by AFOSR grant FA9550-11-1-0016, MURI under AFOSR grant FA9550-10-1-0573, NSF Grant CCF-1065352, and NSF grant CCF-0939370.

C. Quinn is with the Department of Electrical and Computer Engineering, University of Illinois, Urbana, Illinois 61801 quinn7@illinois.edu

N. Kiyavash is with the Department of Industrial and Enterprise Systems Engineering, University of Illinois, Urbana, Illinois 61801 kiyavash@illinois.edu

T. P. Coleman is with the Department of Bioengineering, University of California, San Diego, La Jolla, CA 92093 tpcoleman@ucsd.edu




Fig. 1. A plot of simulated corn and chicken prices. The price of chicken depends on the cost of corn, but not vice versa. Understanding this causal relationship could lead to better investment strategies.




Fig. 2. The traffic influence structure of a computer network during an attack. Unlike the connectivity structure alone, the causal influence structure clearly reveals the attacker.

(see Fig 4).

As illustrated in these examples, the structure of networks of dynamic stochastic systems is not fully captured by considering mutual relationships. Thus, while criteria such as mutual information are used to identify relationships in static stochastic systems, another criterion is needed to define the minimal description of the relationship between the past of some processes and the future of others.

Clive Granger, a Nobel laureate, proposed a framework for this purpose. He developed a methodology for deciding when, in a statistical sense, one process X causally influences another process Y in a network. His methodology is based on quantifying how much causal side information of X , its past, helps in sequentially predicting the future of Y . His framework also accounts for other processes influencing Y , thus characterizing which influences are direct. Researchers have applied his framework in a number of fields, including biology, economics, and social sciences [7], [8], [9], [10], [11], [12]. For most applications, however, the analysis was restricted to the context of linear multivariate models. As we will show, in a general formulation of the sequential prediction task with causal side information, an information theoretic quantity known as causally conditioned directed information [13] is precisely the value of having access to the causal side




Fig. 3. To understand how trends, news, and gossip transfer dynamically in a social network, the relationship structure needs to be understood, not just the connectivity structure.




Fig. 4. Visual processing in the brain occurs with information propagating between various parts of the brain. In addition to feed-forward pathways from the eye to higher-level processing regions of the brain, feedback has been shown to play a crucial role. The feedback might not be evident from only considering purely correlative relationships.²

information.

We will thus propose a graphical model based on causally conditioned directed information. Each node will represent a random process. An edge will be drawn from a node X to a node Y if X influences Y in the sense of Granger, quantified by causally conditioned directed information. The causal conditioning is on all other processes to account for indirect influences. The resulting graphs, called directed information graphs, are analogous to Markov networks in terms of Markov blankets, graph separation, and I-map properties.

We further justify that directed information graphs represent the causal influence structure by showing that they are equivalent to graphs based on generative models. For the latter, the joint distribution is factorized over time, analogous to a state-space model of coupled differential equations. The terms in the factorization are then reduced. These minimal generative models are analogous to Bayesian networks. Interestingly, while Markov and Bayesian networks are distinct, offering different perspectives on mutual dependence structure between variables, directed information graphs and minimum generative model graphs offer the same fundamental perspective on the causal influence structure between processes.

Since the directed information graphs and minimum generative model graphs are equivalent, either can be used to identify the structure of the system. Despite this flexibility, finding the structure by either definition requires computing quantities

²Figure adapted from http://commons.wikimedia.org/wiki/File:Human_Brain_sketch_with_eyes_and_cerebrillum.svg

using the full joint distribution. A natural question is whether instead of the full joint distribution, lower-dimensional statistics could be used when there is some knowledge about the structure. For instance, some gene networks, ecosystems, and computer virus infection networks have random graph structures, and some metabolic, neuronal, and social networks have small-world structures [14]. Optimal transportation systems and the blood vessel system are known to have tree structures [15], [16].

Trees are among the simplest of structures, and can be found using only pairwise statistics between processes. Chow and Liu proposed a method for finding the best tree structured Bayesian network approximation for a set of random variables [17]. Recently, analogous approaches for recovering tree structures in context of networks of random processes were suggested in [18] and [19]. These approaches use pairwise statistics and a coupled, global optimization step. We will propose a method that recovers the structure using the minimal-dimension statistics necessary when the in-degrees of the processes are upper bounded. Also, the method will not require any coupled optimization step.

A. Our Contribution

We propose two graphical models for identifying the structure of networks of causally interacting, stochastic, dynamic processes. The first, minimum generative model graphs, is based on reduced factorizations of the joint distribution of the network over time. It is motivated by simplifying sets of coupled differential equations. The other, directed information graphs, is motivated by Granger causality. We show that directed information quantifies Granger causality in a general prediction framework. Causal influences between each pair of processes is directly queried using directed information. Moreover, we show that these two graphical models are analogous to Markov and Bayesian networks, respectively. Unlike Markov and Bayesian networks, however, the two proposed graphical models are equivalent, suggesting they encode a fundamental causal influence structure of the network.

We discuss how the structure of directed information graphs is related to causal Markov chains. We show that the variable dependence structure underneath a directed information graph is a dynamic Bayesian network. Additionally, we demonstrate that directed information graphs, similar to Markov and Bayesian networks, form a type of independence map (I-map) with a graphical separation criterion similar to d-separation for Bayesian networks.

Finally, we propose an efficient method to identify the structure of directed information graphs when the in-degrees are bounded. The method uses the minimal-dimension statistics necessary to find the structure. For a process \mathbf{Y} with at most K parents, the method uses only $(K + 1)$ -wise statistics. Furthermore, the method finds the parents of \mathbf{Y} independently of finding the parents of other processes. We also demonstrate this method by inferring the minimal generative model structure of a network of causally interacting processes.

B. Related Work

1) *Directed information*: Directed information was first introduced by Marko [20] and independently rediscovered by Rissanen and Wax [21]. Rissanen and Wax proposed their work as an extension of Granger's framework [22], which in turn was based on Wiener's work [23]. Directed information was later formalized by Massey [24]. It plays a fundamental role in communication with feedback [20], [13], [25], [26], [24], [27], gambling with causal side information [28], [29], control over noisy channels [30], [31], [32], [25], and source coding with feed forward [33], [29].

Directed information has already been used in some applications to infer causal relationships between nonlinear processes. In Marko's original paper [20], directed information is used to analyze social relationships between primates. Other applications of directed information include analysis of neuroscience data [34], [35], [36], [37], gene regulatory data [38], [39], and video recordings [40]. Parametric methods have been proposed for estimating the directed information in context of point processes [34], [35], [37]. Additionally, a universal estimation technique applicable to discrete time, finite alphabet processes was proposed in [41].

2) *Graphical models and Granger causality*: There is a large body of literature on graphical models for representing the mutual dependence structure of sets of random variables. We will only reference Pearl [42] and Koller and Friedman [5], which provide good foundations for the theory and techniques. These models have been applied to networks of causally interacting processes. One approach is dynamic Bayesian networks, where each variable for each process is represented as a node in the graph [43]. However, these graphs unravel to larger and larger size with time, depend upon a spatial ordering of the processes using the chain rule, and do not succinctly elucidate relationships between the past of some processes and the future of others.

In this paper, we use the the framework of Granger causality. An alternative framework for identifying causal interactions between variables is based on the principle of intervention [44]. By fixing certain variables to have specific values, observing how the statistics of other variables change can be used to infer some causal relationships. Ay and Polani [45] propose a measure for how strong such causal effects are in the context of processes. Some connections between Pearl's framework and Granger causality (as measured by directed information) are explored in [46].

There are a class of graphical models developed to represent Granger's principle, known as Granger causality graphs [47], [48], [49]. These are mixed graphs (both directed and undirected graphs) for multivariate autoregressive time series. Nodes represent processes. The directed edges represent causal influences, as measured by Granger causality. The undirected edges represent instantaneous correlation.

In [47], it is suggested that, conceptually, Granger causality graphs could be employed for nonlinear relationships. However, it is mentioned that some properties of the graphical model would not hold. They also suggested it would be impossible to infer structures where causal influences were nonlinear without assuming specific models.

There was a recent work on Granger causality graphs that proposed using directed information [50]. However, it justifies using directed information conceptually, motivated by equivalence of Granger causality and directed information in the case of jointly Gaussian processes [51], and does not identify properties of the graph. Independently, [52] showed the relationship between Granger causality and transfer entropy. Directed information is the time average of transfer entropy. Transfer entropy was proposed by Schreiber [53], independently of directed information.

Other works proposing the use of directed information to quantify Granger causality for networks of general processes include [34], [35], [36], [40], independent of each other and of [50], [51]. The work [34] also proposed a graphical model, though different than what is proposed here.

3) *Graph structure identification*: When an upper bound on the in-degree is known for the structure of a set of random variables, there are algorithms for Markov and Bayesian networks that use local statistics (from pairwise up to the size of the upper bound). They recover the Markov blankets and consequently the structure. An early work using this approach is the SGS algorithm [54]. More recent works are discussed in [55] and [5]. An alternative approach to identifying sparse structures of networks of causal processes uses group Lasso, based on the model selection technique Lasso [56].

C. Paper Organization

In Section II we establish definitions and notations. We also review basic properties of Bayesian and Markov networks that will be relevant for our proposed graphical models. In Section III we formally introduce the two proposed graphical models. We discuss how directed information precisely quantifies Granger causality in a general, sequential prediction setting. We then show the two proposed graphical models are equivalent. In Section IV, we discuss how the variable structure underlying the proposed graphical models is the unique, causal Bayesian network. We also formulate how causally conditioned independences are represented in directed information graphs by describing their I-map properties. In Section V we consider algorithms for identifying the graphical models. We describe the procedure that identifies the directed information graph with minimal-dimension statistics when upper bounds on the in-degrees are known. In Section VI, we use this efficient method to infer the structure of a simulated network of causally interacting processes. In the appendices, we provide the proofs of the lemmas and theorems in the paper.

II. BACKGROUND

A. Notation and Information Theoretic Definitions

- For a sequence a_1, a_2, \dots , denote (a_i, \dots, a_j) as a_i^j and $a^k \triangleq a_1^k$.
- Denote $[m] \triangleq \{1, \dots, m\}$ and the power set $2^{[m]}$ on $[m]$ to be the set of all subsets of $[m]$.
- For any Borel space Z , denote its Borel sets by $\mathcal{B}(Z)$ and the space of probability measures on $(Z, \mathcal{B}(Z))$ as $\mathcal{P}(Z)$.

- Consider two probability measures \mathbb{P} and \mathbb{Q} on $\mathcal{P}(Z)$. \mathbb{P} is absolutely continuous with respect to \mathbb{Q} ($\mathbb{P} \ll \mathbb{Q}$) if $\mathbb{Q}(A) = 0$ implies that $\mathbb{P}(A) = 0$ for all $A \in \mathcal{B}(Z)$. If $\mathbb{P} \ll \mathbb{Q}$, denote the Radon-Nikodym derivative as the random variable $\frac{d\mathbb{P}}{d\mathbb{Q}} : Z \rightarrow \mathbb{R}$ that satisfies

$$\mathbb{P}(A) = \int_{z \in A} \frac{d\mathbb{P}}{d\mathbb{Q}}(z) \mathbb{Q}(dz), \quad A \in \mathcal{B}(Z).$$

- The *Kullback-Leibler divergence* between $\mathbb{P} \in \mathcal{P}(Z)$ and $\mathbb{Q} \in \mathcal{P}(Z)$ is defined as

$$D(\mathbb{P} \parallel \mathbb{Q}) \triangleq \mathbb{E}_{\mathbb{P}} \left[\log \frac{d\mathbb{P}}{d\mathbb{Q}} \right] = \int_{z \in Z} \log \frac{d\mathbb{P}}{d\mathbb{Q}}(z) \mathbb{P}(dz)$$

if $\mathbb{P} \ll \mathbb{Q}$ and ∞ otherwise.

- Throughout this paper, we will consider m random processes where the i th (with $i \in \{1, \dots, m\}$) random process at time j (with $j \in \{1, \dots, n\}$), takes values in a Borel space X .
- For a sample space Ω , sigma-algebra \mathcal{F} , and probability measure \mathbb{P} , denote the probability space as $(\Omega, \mathcal{F}, \mathbb{P})$.
- Denote the i th random variable at time j by $X_{i,j} : \Omega \rightarrow X$, the i th random process as $\mathbf{X}_i = (X_{i,1}, \dots, X_{i,n})$, the subset of random processes $\underline{\mathbf{X}}_{\mathcal{I}} = (\mathbf{X}_i : i \in \mathcal{I})$, and the whole collection of all m random processes as $\underline{\mathbf{X}} \triangleq \underline{\mathbf{X}}_{[m]}$. Denote the whole collection of all m random processes from time j to j' as $\underline{\mathbf{X}}_j^{j'}$.
- The probability measure \mathbb{P} thus induces a probability distribution on $X_{i,j}$ given by $P_{X_{i,j}}(\cdot) \in \mathcal{P}(X)$, a joint distribution on \mathbf{X}_i given by $P_{\mathbf{X}_i}(\cdot) \in \mathcal{P}(X^n)$, and a joint distribution on $\underline{\mathbf{X}}_{\mathcal{I}}$ given by $P_{\underline{\mathbf{X}}_{\mathcal{I}}}(\cdot) \in \mathcal{P}(X^{|\mathcal{I}|n})$.
- With slight abuse of notation, denote $\mathbf{Y} \equiv \mathbf{X}_i$ for some i and $\mathbf{X} \equiv \mathbf{X}_k$ for some $i \neq k$ and denote the conditional distribution and *causally conditioned* distribution of \mathbf{Y} given \mathbf{X} respectively as

$$\begin{aligned} P_{\mathbf{Y}|\mathbf{X}=\mathbf{x}}(d\mathbf{y}) &\triangleq P_{\mathbf{Y}|\mathbf{X}}(d\mathbf{y}|\mathbf{x}) \\ &= \prod_{j=1}^n P_{Y_j|Y^{j-1}, X^n}(dy_j|y^{j-1}, x^n) \end{aligned} \quad (1)$$

$$\begin{aligned} P_{\mathbf{Y}||\mathbf{X}=\mathbf{x}}(d\mathbf{y}) &\triangleq P_{\mathbf{Y}||\mathbf{X}}(d\mathbf{y}||\mathbf{x}) \\ &\triangleq \prod_{j=1}^n P_{Y_j|Y^{j-1}, X^{j-1}}(dy_j|y^{j-1}, x^{j-1}). \end{aligned} \quad (2)$$

Note the similarity between regular conditioning (1) and (2), except in causal conditioning the future (x_j^n) is not conditioned on (Ch. 3 of [13])³.

- With slight abuse of notation, $\underline{\mathbf{W}} \equiv \underline{\mathbf{X}}_{\mathcal{I}}$ for some $\mathcal{I} \subseteq [m] \setminus \{i\}$ with $\mathcal{W} = X^{|\mathcal{I}|n}$. Consider two sets of causally conditioned distributions $\{P_{\mathbf{Y}||\underline{\mathbf{W}}=\underline{\mathbf{w}}} \in \mathcal{P}(\mathbf{Y}) : \underline{\mathbf{w}} \in \mathcal{W}\}$ and $\{Q_{\mathbf{Y}||\underline{\mathbf{W}}=\underline{\mathbf{w}}} \in \mathcal{P}(\mathbf{Y}) : \underline{\mathbf{w}} \in \mathcal{W}\}$ along with a marginal distribution $P_{\underline{\mathbf{W}}} \in \mathcal{P}(\mathcal{W})$. Then the conditional

³Note the slight difference in conditioning upon x^{j-1} in this definition as compared to x^j in the original causal conditioning definition. The purpose for doing this will become clear later in the manuscript.

KL divergence is given by

$$\begin{aligned} D(P_{Y||\underline{W}}||Q_{Y||\underline{W}}|P_{\underline{W}}) \\ = \int_{\underline{W}} D(P_{Y||\underline{W}=\underline{w}}||Q_{Y||\underline{W}=\underline{w}}) P_{\underline{W}}(d\underline{w}). \end{aligned}$$

The following Lemma will be useful throughout:

Lemma 2.1: $D(P_{Y||\underline{W}}||Q_{Y||\underline{W}}|P_{\underline{W}}) = 0$ if and only if $P_{Y||\underline{W}=\underline{w}}(dy) = Q_{Y||\underline{W}=\underline{w}}(dy)$ with $P_{\underline{W}}$ probability one.

- Let $\mathbf{X} \equiv \mathbf{X}_i$ for some i , $\mathbf{Y} \equiv \mathbf{X}_k$ for some k and $\underline{\mathbf{W}} \equiv \underline{\mathbf{X}}_{\mathcal{I}}$ for some $\mathcal{I} \subseteq [m] \setminus \{i, k\}$. The mutual information, *directed information* [24], and causally conditioned directed information [13] are given by

$$\begin{aligned} I(\mathbf{X}; \mathbf{Y}) &\triangleq D(P_{Y|\mathbf{X}}||P_Y|P_{\mathbf{X}}) \\ I(\mathbf{X} \rightarrow \mathbf{Y}) &\triangleq D(P_{Y|\mathbf{X}}||P_Y|P_{\mathbf{X}}) \end{aligned} \quad (3)$$

$$I(\mathbf{X} \rightarrow \mathbf{Y}||\underline{\mathbf{W}}) \triangleq D(P_{Y|\mathbf{X},\underline{\mathbf{W}}}|P_{Y||\underline{\mathbf{W}}}|P_{\mathbf{X},\underline{\mathbf{W}}}). \quad (4)$$

Conceptually, mutual information and directed information are related. However, while mutual information quantifies statistical correlation (in the colloquial sense of statistical interdependence), directed information quantifies statistical *causation*. We later justify this statement showing that directed information is a general formulation of Granger causality. Note that $I(\mathbf{X}; \mathbf{Y}) = I(\mathbf{Y}; \mathbf{X})$, but $I(\mathbf{X} \rightarrow \mathbf{Y}) \neq I(\mathbf{Y} \rightarrow \mathbf{X})$ in general. As a consequence of Lemma 2.1 and (4), we have:

Corollary 2.2: $I(\mathbf{X} \rightarrow \mathbf{Y}||\underline{\mathbf{W}}) = 0$ if and only if \mathbf{Y} is causally independent of \mathbf{X} causally conditioned on $\underline{\mathbf{W}}$:

$$P_{Y||\mathbf{X}=\mathbf{x},\underline{\mathbf{W}}=\underline{\mathbf{w}}}(dy) = P_{Y||\underline{\mathbf{W}}=\underline{\mathbf{w}}}(dy), \quad P_{\mathbf{X},\underline{\mathbf{W}}} - a.s.$$

Equivalently, we denote that $\mathbf{X} \rightarrow \underline{\mathbf{W}} \rightarrow \mathbf{Y}$ forms a *causal Markov chain*. This is analogous to Markov chains, denoted as $\mathbf{X} - \underline{\mathbf{W}} - \mathbf{Y}$, where $I(\mathbf{X}; \mathbf{Y}||\underline{\mathbf{W}}) = 0$ if and only if \mathbf{Y} is independent of \mathbf{X} conditioned on $\underline{\mathbf{W}}$:

$$P_{Y|\mathbf{X}=\mathbf{x},\underline{\mathbf{W}}=\underline{\mathbf{w}}}(dy) = P_{Y||\underline{\mathbf{W}}=\underline{\mathbf{w}}}(dy), \quad P_{\mathbf{X},\underline{\mathbf{W}}} - a.s.$$

- Let $G = (V, E)$ denote a directed graph. For each edge $(u, v) \in E$, u is called the *parent* and v is the *child*.

We next review basic properties of Markov and Bayesian networks, as the proposed graphical models are analogous to these. The remaining discussion in this section follows from Chapter 3 of [42].

B. Bayesian Networks

To motivate Bayesian networks, consider the following example.

Example 1: Let $\{W, X, Y, Z\}$ be a set of four random variables with a positive joint distribution $P_{W,X,Y,Z}$. Let their relationships be of the form:

$$\begin{aligned} Y &= W + X + \xi \\ Z &= W + \xi' \end{aligned}$$

where ξ and ξ' are noises and W , X , ξ , and ξ' are all independent. Consequently, the joint distribution can be factorized




Fig. 5. Bayesian and Markov networks for Example 1.

as:

$$\begin{aligned} P_{W,X,Y,Z}(w, x, y, z) \\ = P_W(w)P_{X|W}(x|w)P_{Y|W,X}(y|w, x)P_{Z|W,X,Y}(z|w, x, y) \\ = P_W(w)P_X(x)P_{Y|W,X}(y|w, x)P_{Z|W}(z|w) \end{aligned}$$

Since the distribution cannot be reduced further, these dependencies represent the structure. This structure can be depicted graphically with directed edges corresponding to dependencies, as is shown in Figure 5(a).

Figure 5(a) is an example of a Bayesian network. Bayesian networks are directed graphs representing conditional dependencies in a reduced factorization of the joint distribution. Note that the Bayesian network in Example 1 depended on how the chain rule was applied to the joint distribution. Other orderings correspond to different graphs. For instance, the chain rule could be applied as $P_{W,X,Y,Z}(w, x, y, z) = P_Y(y)P_{Z|Y}(z|y)P_{X|Y,Z}(x|y, z)P_{W|X,Y,Z}(w|x, y, z)$. Here, however, no term can be reduced.

Also note that there are no directed cycles. This is because a variable can only have incoming arrows from variables with smaller index and outgoing arrows to variables with a larger index. Lastly, although the graph has directed edges, the edges correspond to conditional dependence relationships, which are mutual. We now discuss an alternative representation.

C. Markov Networks

Bayesian networks are one method to graphically represent the conditional dependence structure of a set of random variables. It is based on factorizations of the joint distribution, where unnecessary dependencies are removed. An alternative method is known as Markov networks, which are undirected graphs. An edge is drawn between each pair of variables that are dependent, given knowledge of all other variables. Conditional mutual information is used to quantify dependence.

Consider the system in Example 1 (see Figure 5). For the Markov network, each edge is determined by testing “globally” conditioned dependencies. For instance,

$$\begin{aligned} Y &\perp\!\!\!\perp Z \mid W, X \\ X &\perp\!\!\!\perp Z \mid W, Y \\ W &\not\perp\!\!\!\perp X \mid Y, Z. \end{aligned}$$

The remaining three pairs are conditionally dependent. Note that the last formula above, for instance, corresponds to

$I(W; X|Y, Z) > 0$. The Markov network for Example 1 is shown in Figure 5(b).

Markov networks are undirected graphs and do not depend on variable orderings. Also, as shown in Figure 5(b), Markov networks can have loops. However, note that W and X share an edge in the Markov network. Even though they are marginally independent, Y depends on both. Thus, conditioning on Y induces dependence between W and X .

The Bayesian network in Example 1 did not have an edge between W and X , although some other variable orderings would result in such an edge. Although describing the same set of random variables, Markov and Bayesian networks represent different mutual dependencies. We next consider how to succinctly represent how the past of some processes statistically affect the future of others in a network of random processes.

III. MINIMAL GENERATIVE MODEL GRAPHS AND DIRECTED INFORMATION GRAPHS

We now consider the specific case of physically causal, stochastic dynamical systems. Here, there are multiple processes, each of which is a indexed collection of random variables, and the indexing corresponds to time. Our goal is to graphically represent the structure of causal influences between processes. We consider two approaches.

We first introduce a graphical model based on generative models for stochastic dynamical systems. It is motivated by the process of reducing dependencies in coupled differential equations for deterministic dynamic systems. The other graphical model is motivated by the framework of Granger causality, which directly tests for causal influence. Causally conditioned directed information is shown to be a general formulation of Granger causality and used as the edge selection criterion. Thus, there is an analogy between the graphical models proposed in this section and Bayesian and Markov models respectively. Although Bayesian and Markov models are different representations of the structure, we will lastly show that despite different motivations and methodologies, the two proposed graphical models are identical. These graphical models were first discussed in preliminary work in [57].

A. Minimal Generative Model Graphs

Stochastic dynamical systems have indexing both across processes and across *time*. They have a natural formula representation - the coupled differential equations that characterize the dynamics of the system over time. Consider the following example of a simple system:

Example 2: Let x_t , y_t , and z_t be three processes comprising a physical, dynamical system. The evolution of the processes over time can be fully described by coupled differential equations:

$$\dot{x} = f(x, y, z) \quad (5a)$$

$$\dot{y} = g(x, y, z) \quad (5b)$$

$$\dot{z} = h(x, y, z). \quad (5c)$$




Fig. 6. A graphical model of the causal influence structure of the stochastic dynamical system of Example 2.

For small Δ , this becomes

$$x_{t+\Delta} = x_t + \Delta f(x_t, y_t, z_t) \quad (6a)$$

$$y_{t+\Delta} = y_t + \Delta g(x_t, y_t, z_t) \quad (6b)$$

$$z_{t+\Delta} = z_t + \Delta h(x_t, y_t, z_t). \quad (6c)$$

Note that the dynamics in this system are causal. Given the full past of the system, $x_{t+\Delta}$ can be completely generated without knowledge of $y_{t+\Delta}$ or $z_{t+\Delta}$. Thus, although these equations are coupled, knowledge of the full past decouples the dynamics of the near future. Suppose that the processes are not each dependent on the past of *all* other processes. For instance, suppose x_t and y_t influence each other independently of z_t , and that z_t does not depend on x_t :

$$\dot{x} = f(x, y,)$$

$$\dot{y} = g(x, y,)$$

$$\dot{z} = h(, y, z).$$

For small Δ , this becomes

$$x_{t+\Delta} = x_t + \Delta f(x_t, y_t,)$$

$$y_{t+\Delta} = y_t + \Delta g(x_t, y_t,)$$

$$z_{t+\Delta} = z_t + \Delta h(, y_t, z_t).$$

In this scenario, we can construct a directed graph where, for each node, we place incoming edges from those nodes that describe the minimal set of processes whose past affect the near future of process pertaining to the node of interest. The physically causal dependencies represented in the reduced formulas characterizes the structure of this dynamical system. See Figure 6.

Now consider the system analogous to (5) but where $(X, Y, Z) = (X_t, Y_t, Z_t : t \geq 0)$ is a multi-dimensional Ito process satisfying the stochastic differential equations [58]:

$$dX = f(X_t, Y_t, Z_t)dt + dB_t \quad (8a)$$

$$dY = g(X_t, Y_t, Z_t)dt + dC_t \quad (8b)$$

$$dZ = h(X_t, Y_t, Z_t)dt + dD_t \quad (8c)$$

where $(B, C, D) = (B_t, C_t, D_t : t \geq 0)$ are independent Wiener processes. An analogous relation holds for a multi-dimensional jump processes where $(B, C, D) = (B_t, C_t, D_t : t \geq 0)$ are independent Poisson processes.

For sufficiently small Δ , this becomes:

$$X_{t+\Delta} = X_t + \Delta f(X_t, Y_t, Z_t) + (B_{t+\Delta} - B_t)$$

$$Y_{t+\Delta} = Y_t + \Delta g(X_t, Y_t, Z_t) + (C_{t+\Delta} - C_t)$$

$$Z_{t+\Delta} = Z_t + \Delta h(X_t, Y_t, Z_t) + (D_{t+\Delta} - D_t).$$

Note that due to the independent increments property of the Wiener and Poisson process, $(B_{t+\Delta} - B_t)$ is statistically independent of $(X_\tau, Y_\tau, Z_\tau : \tau \leq t)$ and is also statistically independent of $(C_{t+\Delta} - C_t)$ and $(D_{t+\Delta} - D_t)$.

Thus, analogous to (6), knowledge of the full past statistically decouples the dynamics of the multi-dimensional process in the near future.

If we discretize time into units of length Δ , then the stochastic dynamics are fully captured by the joint distribution P_{X^n, Y^n, Z^n} . Analogous to the coupled differential equations, the chain rule can be applied over time to obtain a similar representation:

$$P_{X^n, Y^n, Z^n}(dx^n, dy^n, dz^n) = \prod_{j=1}^n P_{X_j, Y_j, Z_j | X^{j-1}, Y^{j-1}, Z^{j-1}}(dx_j, dy_j, dz_j | x^{j-1}, y^{j-1}, z^{j-1}). \quad (9)$$

Due to the dynamics of the multi-dimensional process in the near future being statistically decoupled given knowledge of the full past, we have:

$$P_{X^n, Y^n, Z^n}(dx^n, dy^n, dz^n) = \prod_{j=1}^n P_{X_j | X^{j-1}, Y^{j-1}, Z^{j-1}}(dx_j | x^{j-1}, y^{j-1}, z^{j-1}) \times P_{Y_j | X^{j-1}, Y^{j-1}, Z^{j-1}}(dy_j | x^{j-1}, y^{j-1}, z^{j-1}) \times P_{Z_j | X^{j-1}, Y^{j-1}, Z^{j-1}}(dz_j | x^{j-1}, y^{j-1}, z^{j-1}). \quad (10)$$

Now suppose we have a multi-dimensional stochastic differential equation of the form:

$$\begin{aligned} dX &= f(X_t, Y_t)dt + dB_t \\ dY &= g(X_t, Y_t)dt + dC_t \\ dZ &= h(Y_t, Z_t)dt + dD_t. \end{aligned}$$

For sufficiently small Δ , this becomes:

$$\begin{aligned} X_{t+\Delta} &= X_t + \Delta f(X_t, Y_t) + (B_{t+\Delta} - B_t) \\ Y_{t+\Delta} &= Y_t + \Delta g(X_t, Y_t) + (C_{t+\Delta} - C_t) \\ Z_{t+\Delta} &= Z_t + \Delta h(Y_t, Z_t) + (D_{t+\Delta} - D_t). \end{aligned}$$

With this, we can further remove the unnecessary dependencies in the discrete-time model to obtain:

$$P_{X^n, Y^n, Z^n}(dx^n, dy^n, dz^n) = \prod_{j=1}^n P_{X_j | X^{j-1}, Y^{j-1}}(dx_j | x^{j-1}, y^{j-1}) \times P_{Y_j | X^{j-1}, Y^{j-1}}(dy_j | x^{j-1}, y^{j-1}) \times P_{Z_j | Y^{j-1}, Z^{j-1}}(dz_j | y^{j-1}, z^{j-1}). \quad (12)$$

This causal dependence structure is shown graphically in Figure 6.

We can extend this stochastic differential equation relation to the following lemma:

Lemma 3.1: Let B be an m -dimensional Brownian motion and denote \mathcal{F}_t as the sigma-algebra generated by $(B_\tau, \tau < t)$. Assume that $u = [u_j(t, \omega) : t \in [0, T], 1 \leq j \leq m]$ and $v = [v_{i,j}(t, \omega) : t \in [0, T], 1 \leq i, j \leq m]$ have the property that u_t and v_t are \mathcal{F}_t -measurable for all $t \in [0, T]$. Assume

that $v_{i,j}(t, \omega) = 0$ for $i \neq j$. Then the m -dimensional Ito process $X = [X_j(t, \omega) : t \in [0, T], 1 \leq j \leq m]$ satisfying the stochastic differential equation of the form

$$dX_t = udt + vdB_t$$

has the conditional independence property

$$\mathbb{P}(X_t \in A_1 \times \cdots \times A_m | \mathcal{F}_t) = \prod_{i=1}^m \mathbb{P}(X_{i,t} \in A_i | \mathcal{F}_t).$$

Proof: Note that the integral equation becomes [58, Ch 4]

$$X_t = X_0 + \int_0^t u(s, \omega)ds + \int_0^t v(s, \omega)dB_s.$$

Also note that we can pick the left endpoint $v_{t_{i-1}}$ in the Ito definition [58, Ch 4] of $\int_0^t v(s, \omega)dB_s$, so the latter becomes

$$\lim_{n \rightarrow \infty} \sum_{[t_{i-1}, t_i] \in \mathcal{P}_n} v_{t_{i-1}}(B_{t_i} - B_{t_{i-1}})$$

where $\{\mathcal{P}_n\}$ is a sequence of partitions of $[0, T]$. The proof then follows directly by the assumed diagonal structure of v , the fact that an m dimensional Wiener process is a set of m independent Wiener processes, and the independent increments property of the Wiener process. ■

Definition 3.2: A distribution $P_{\underline{\mathbf{X}}}$ is called *positive* if there exists a measure ϕ such that $P_{\underline{\mathbf{X}}} \ll \phi$ and $\frac{dP_{\underline{\mathbf{X}}}}{d\phi}(\underline{\mathbf{x}}) > 0$ for all $\underline{\mathbf{x}}$ in the support of $P_{\underline{\mathbf{X}}}$.

Assumption 1: For the remainder of this paper, we only consider joint distributions that are positive and that satisfy (10).

Remark 1: The assumption of positivity is to avoid degenerate cases that arise with deterministic relationships. For instance, suppose X is a random process with a continuous distribution and Y represents X passed through a deterministic invertible system. Then X and Y have joint distribution that is not asymptotically continuous with respect to the Lebesgue measure because the distribution of Y given X is a point mass and vice versa. The assumption (10) is implicit in Marko's Kirchhoff current laws pertaining to directed information [20]. Specifically, the assumption (10) must hold in order to equate [20, eqn 14] with [20, eqn 13]. Moreover, the assumption of (10) holds for any continuous-time generative model of a stochastic differential equation satisfying (8) where B, C, D are Wiener or Poisson processes. As such, these are rather modest assumptions that describe models of many important physical systems.

We now generalize this process. Consider a causal, stochastic dynamical system of m processes with a positive joint distribution. Let $\underline{\mathbf{X}}$ denote the set of random processes and let $P_{\underline{\mathbf{X}}}$ denote their joint distribution. The dynamics of the system is fully described by $P_{\underline{\mathbf{X}}}$. $P_{\underline{\mathbf{X}}}$ could initially be factorized over space (index of the processes) or time. Like in (9), we first apply the chain rule over time:

$$P_{\underline{\mathbf{X}}}(d\underline{\mathbf{x}}) = \prod_{j=1}^n P_{\underline{\mathbf{X}}_j | \underline{\mathbf{X}}^{j-1}}(d\underline{\mathbf{x}}_j | \underline{\mathbf{x}}^{j-1}).$$

With Assumption 1, given the full past, the dynamics of the processes in the near future decouple analogous to (10):

$$P_{\underline{\mathbf{X}}}(d\underline{\mathbf{x}}) = \prod_{j=1}^n \prod_{i=1}^m P_{X_{i,j}|\underline{\mathbf{x}}^{j-1}}(dx_{i,j}|\underline{\mathbf{x}}^{j-1}). \quad (13)$$

Equivalently, using causal conditioning notation (2):

$$P_{\underline{\mathbf{X}}}(d\underline{\mathbf{x}}) = \prod_{i=1}^m P_{\mathbf{X}_i|\underline{\mathbf{x}}_{[m]\setminus\{i\}}}(d\mathbf{x}_i|\underline{\mathbf{x}}_{[m]\setminus\{i\}}).$$

By factorizing over time first, each \mathbf{X}_i is still conditioned on the past of every other process. Like in (12), we now remove unnecessary causal dependencies. For each process \mathbf{X}_i , let $A(i) \subseteq [m] \setminus \{i\}$ denote a subset of other processes. Define the corresponding induced probability measure P_A :

$$P_A(d\underline{\mathbf{x}}) = \prod_{i=1}^m P_{\mathbf{X}_i|\underline{\mathbf{x}}_{A(i)}}(d\mathbf{x}_i|\underline{\mathbf{x}}_{A(i)}). \quad (14)$$

We want to pick the subsets $\{A(i)\}_{i=1}^m$ so that their cardinalities are small, while still capturing the full dynamics:

$$D(P_{\underline{\mathbf{X}}}||P_A) = 0. \quad (15)$$

In Example 2, the $A(i)$'s would correspond to $\{\mathbf{Y}\}$, $\{\mathbf{X}\}$, and $\{\mathbf{Y}\}$ for \mathbf{X} , \mathbf{Y} , and \mathbf{Z} , respectively.

Definition 3.3: Under Assumption 1, for a joint distribution $P_{\underline{\mathbf{X}}}$, a *minimal generative model* is a function $A : [m] \rightarrow 2^{[m]}$ where the cardinalities $\{|A(i)|\}_{i=1}^m$ are minimal such that (15) holds.

Note that the $A(i)$'s together must satisfy (15), suggesting that choice of $A(i)$ for a particular i might be restricted by what was already chosen for $\{A(j)\}_{j=1}^{i-1}$. However, by non-negativity of KL divergence, (15) corresponds to

$$D\left(P_{\mathbf{X}_i|\underline{\mathbf{x}}_{[m]\setminus\{i\}}}||P_{\mathbf{X}_i|\underline{\mathbf{x}}_{A(i)}}|P_{\underline{\mathbf{x}}_{[m]\setminus\{i\}}}\right) = 0 \quad (16)$$

for all $i \in [m]$. Thus, the sets can be chosen separately to satisfy this condition. Furthermore, under our assumption, minimum generative models are unique.

Lemma 3.4: For any distribution $P_{\underline{\mathbf{X}}}$ satisfying Assumption 1, the minimal generative model is unique.

The proof appears in Appendix A. Bayesian networks depend on the indexing of the variables. For minimum generative models, however, by first factorizing over time, the indexing of the processes becomes irrelevant. We now define a corresponding graphical model.

Definition 3.5: A *minimum generative model graph* is a directed graph for a minimum generative model, where each process is represented by a node, and there is a directed edge from \mathbf{X}_k to \mathbf{X}_i for $i, k \in [m]$ iff $k \in A(i)$.

Note that unlike Bayesian networks, minimum generative model graphs can have directed loops, as is the case in Figure 6.

Minimum generative model graphs represent reduced factorizations of the joint distribution of the system. They encode causal relationships by only depicting as parents in the graph those subsets of processes that are necessary and sufficient to describe the full dynamics. This graphical model was motivated by reducing coupled differential equations for deterministic systems. We next propose an alternative graphical

model based on the framework of Granger causality, which directly tests for causal influences between processes. We show that causally conditioned directed information captures Granger's principle and use it as an edge selection criterion.

B. Granger Causality and Directed Information Graphs

In 1969, motivated by earlier work by Norbert Wiener [23], Nobel laureate Clive Granger proposed a framework for identifying when one process statistically “causes” another [22]:

“We say that \mathbf{X} is causing \mathbf{Y} if we are better able to predict [the future of] \mathbf{Y} using all available information than if the information apart from [the past of] \mathbf{X} had been used.”

While this definition is general, its previous formulations have been restricted to specific classes of models, such as autoregressive linear models. Specifically, Granger's setup was as follows [22]. Suppose we have three processes \mathbf{X} , \mathbf{Y} , and \mathbf{Z} . Granger posited the development of two autoregressive linear models:

$$Y_t = \sum_{\tau>0} a_\tau Y_{t-\tau} + \sum_{\tau'>0} b_{\tau'} X_{t-\tau'} + c_{\tau'} Z_{t-\tau'} + E_t \quad (17)$$

$$Y_t = \sum_{\tau>0} \tilde{a}_\tau Y_{t-\tau} + \sum_{\tau'>0} \tilde{c}_{\tau'} Z_{t-\tau'} + \tilde{E}_t. \quad (18)$$

Next, he proposed the development of least squares so that each predictor performs rationally to find the best coefficients for (a, b, c) in (17) and analogously for (\tilde{a}, \tilde{c}) in (18). The performance of the predictors is measured by the variability in E_t and \tilde{E}_t , given by variances σ^2 and $\tilde{\sigma}^2$. Note that in general, after the least squares fit, $\tilde{\sigma}^2 \geq \sigma^2$. In the event that $\tilde{\sigma}^2 = \sigma^2$, removing that past of X had no detrimental effect on predicting the future of Y . As such, Granger suggested the calculation of the following quantity:

$$\log \frac{\tilde{\sigma}^2}{\sigma^2}$$

and argued that X causes Y if and only if it is deemed positive (in a statistically significant sense). In many scenarios, the modality of the process Y is not even continuous-valued and thus modeling equations like (17) and (18) are not applicable. For example, suppose we would like to understand the causal influence between neural spike trains, or the causal influence between the timings of packets in a computer network [34]. If we discretize time into small enough bins, then Y_i is 0 or 1, depending on whether or not an event occurred in that bin. However, the right-hand side of (17) and (18) are continuous-valued.

1) *A general sequential prediction approach:* We will consider a general formulation of Granger's statement, emphasizing the words “predict” and “better”. Granger's definition is posed in terms of how one probability model predicts better than another. Let $\underline{\mathbf{X}}$ be all the processes simultaneously being observed. Let $Y^n \equiv \mathbf{X}_i$ denote the stochastic process being predicted. Let $X^n \equiv \mathbf{X}_k$ be another process. The setup is that we are attempting to understand if X^n causes Y^n . Knowledge of the past of X^n is a form of causal side information.

Thus, the overall goal is to characterize how much causal side information helps in sequential prediction. We now formally describe a general sequential prediction setup. (See [59] for an overview.)

Denote \mathcal{F}_t to be the sigma-algebra pertaining to information about the past of all processes, and $\tilde{\mathcal{F}}_t$ to be the sigma-algebra pertaining to information about the past of all processes excluding X^n :

$$\begin{aligned}\mathcal{F}_t &= \sigma(X_{l,\tau} : l \in [m] \setminus \{k\}, \tau < t; X_{k,\tau} : \tau < t) \\ \tilde{\mathcal{F}}_t &= \sigma(X_{l,\tau} : l \in [m] \setminus \{k\}, \tau < t).\end{aligned}$$

Remark 2: Note that if $[m] = \{1, 2\}$ and we denote $X^n \equiv \mathbf{X}_1$ and $Y^n \equiv \mathbf{X}_2$, then this reduces to

$$\begin{aligned}\mathcal{F}_t &= \sigma(X^{t-1}, Y^{t-1}) \\ \tilde{\mathcal{F}}_t &= \sigma(Y^{t-1}).\end{aligned}$$

At time t , one predictor has available information about the past of all processes, and specifies a prediction $q_t \in \mathcal{Q}$ about y_t that is \mathcal{F}_t -measurable. The other predictor has all available information apart from the past of X^n and specifies a prediction $\tilde{q}_t \in \mathcal{Q}$ about y_t that is $\tilde{\mathcal{F}}_t$ -measurable. Define the spaces of candidate predictors as:

$$\begin{aligned}\mathcal{A}_t &= \{q : \Omega \rightarrow \mathcal{Q} \text{ s.t. } q \text{ is } \mathcal{F}\text{-measurable}\} \\ \tilde{\mathcal{A}}_t &= \{\tilde{q} : \Omega \rightarrow \mathcal{Q} \text{ s.t. } \tilde{q} \text{ is } \tilde{\mathcal{F}}\text{-measurable}\}.\end{aligned}$$

Subsequently, y_t is revealed, and a loss function $l : \mathcal{Y} \times \mathcal{Q} \rightarrow \mathbb{R}^+$ dictates that for prediction $p \in \mathcal{Q}$ and outcome y_t , a loss of $l(p, y)$ is incurred. Thus, one predictor incurs loss $l(q, y_t)$ and the other incurs $l(\tilde{q}, y_t)$.

Denote the reduction in loss to characterize whether or not the former prediction was better than the second as:

$$r_t = l(\tilde{q}_t, y_t) - l(q_t, y_t).$$

Analogous to how Granger selected least squares fits for the coefficients in (17) and (18), it is quite natural to select rational agents to minimize loss with respect to the two scenarios, in expectation or in a minimax sense. Afterwards, analogously, the reduction in loss can be discussed as a measure of performance in a worst-case or average sense. We have discussed this in preliminary work in [60]. We will next discuss the logarithmic loss and expected cumulative reduction in loss.

2) *The logarithmic loss:* Now we consider the framework of the logarithmic loss. Let \mathcal{Y} be a measurable space and let μ be a measure on that space. Let the space of predictors be the space of probability measures over \mathcal{Y} that form a density with respect to μ :

$$\mathcal{Q} = \{p \in \mathcal{P}(\mathcal{Y}) : p \ll \mu\}.$$

With this, we define the logarithmic loss as:

$$l(q, y) = -\log \frac{dq}{d\mu}(y) \quad q \ll \mu.$$

Example 3: Suppose that $\mathcal{Y} \subset \mathbb{Z}$ we let μ be the counting measure. Then we have that $l(q, y)$ is simply $-\log q(\{y\})$,

where we note that $q(\{y\})$ is the probability mass function evaluated at outcome y .

We assume that there is a joint distribution $P_{\mathbf{X}}$ on all the random processes and that both agents act rationally in selecting their predictor to minimize loss on average. As such, one agent attempts to minimize his expected loss with respect to $P_{\mathbf{X}}$:

$$q_t^* = \arg \min_{q_t \in \mathcal{A}_t} \mathbb{E}_{P_{\mathbf{X}}} [l(q_t, Y_t)] \quad (19)$$

and the other agent does likewise:

$$\tilde{q}_t^* = \arg \min_{\tilde{q}_t \in \tilde{\mathcal{A}}_t} \mathbb{E}_{P_{\mathbf{X}}} [l(\tilde{q}_t, Y_t)]. \quad (20)$$

The expected cumulative reduction in loss for both rational agents is given by:

$$\bar{R}(q^*, \tilde{q}^*) \triangleq \mathbb{E}_{P_{\mathbf{X}}} \left[\sum_{t=1}^n r_t(q_t^*, \tilde{q}_t^*, Y_t) \right].$$

We now state our main lemma, showing that the optimal predictors are the true conditional distributions and that the reduction in expected loss is precisely the causally conditioned directed information.

Lemma 3.6: The optimal solutions to (19) and (20) are given by

$$\begin{aligned}q_t^* &= P_{Y_t | \mathcal{F}_t} \\ \tilde{q}_t^* &= P_{Y_t | \tilde{\mathcal{F}}_t}.\end{aligned}$$

The expected cumulative reduction in loss is given by the causally conditioned directed information:

$$\bar{R}(q^*, \tilde{q}^*) = I(X^n \rightarrow Y^n \| \mathbf{X}_{[m] \setminus \{i, k\}}).$$

Proof: Note that

$$\begin{aligned}q_t^* &= \arg \min_{q_t \in \mathcal{A}_t} \mathbb{E}_{P_{\mathbf{X}}} \left[-\log \frac{dq_t}{d\mu}(Y_t) \right] \\ &= \arg \min_{q_t \in \mathcal{A}_t} \mathbb{E}_{P_{\mathbf{X}}} \left[-\log \frac{dP_{Y_t | \mathcal{F}_t}}{d\mu}(Y_t) + \log \frac{dP_{Y_t | \mathcal{F}_t}}{dq_t} \right] \\ &= \arg \min_{q_t \in \mathcal{A}_t} D(P_{Y_t | \mathcal{F}_t} \| q_t)\end{aligned} \quad (21)$$

where (21) follows from the definition of divergence and that the left hand term in the expectation does not effect the arg min. Moreover, note that clearly $P_{Y_t | \mathcal{F}_t}$ is \mathcal{F}_t -measurable and thus from the non-negativity of KL divergence, $q_t^* = P_{Y_t | \mathcal{F}_t}$. Similarly, $\tilde{q}_t^* = P_{Y_t | \tilde{\mathcal{F}}_t}$.

We now discuss using Granger's notion of "better" to address the two predictors. Since clearly $q_t^* \ll \tilde{q}_t^*$, note that the reference measure μ disappears and the reduction in loss becomes a log-likelihood ratio:

$$r_t(q_t^*, \tilde{q}_t^*, y_t) = \log \frac{dq_t^*}{d\tilde{q}_t^*}(y_t) = \log \frac{dP_{Y_t | \mathcal{F}_t}}{dP_{Y_t | \tilde{\mathcal{F}}_t}}(y_t).$$

Thus,

$$\begin{aligned}\bar{R}(q^*, \tilde{q}^*) &= \mathbb{E}_{P_{\mathbf{X}}} \left[\sum_{t=1}^n r_t(q_t^*, \tilde{q}_t^*, Y_t) \right] \\ &= \mathbb{E}_{P_{\mathbf{X}}} \left[\sum_{t=1}^n \log \frac{dP_{Y_t|\mathcal{F}_t}}{dP_{Y_t|\tilde{\mathcal{F}}_t}}(Y_t) \right] \\ &= I(X^n \rightarrow Y^n \| \underline{\mathbf{X}}_{[m] \setminus \{i,k\}}).\end{aligned}$$

Note that this embodies the notion of “better” because of the non-negativity of directed information. As such, a natural way to generalize Granger’s notion of causality is say that X causes Y if and only if $I(X^n \rightarrow Y^n \| \underline{\mathbf{X}}_{[m] \setminus \{i,k\}}) > 0$.

Remark 3: Note that the formulation of Granger causality for linear autoregressive models is equivalent to directed information when $P_{\mathbf{X}}$ is jointly Gaussian [52], [50]. We note that while directed information does capture Granger causality in a general setting (expected cumulative regret under the logarithmic loss of the best possible sequential predictors), it does not for all contexts. For instance, if the worst-case regret over outcomes is used, then the resulting value of causal side information is different from directed information.

Lemma 3.6 shows that directed information (3) may be used to also quantify how much causal knowledge of process \mathbf{X} helps in sequentially predicting process \mathbf{Y} , in the sense of Granger. We now define a graphical model, using directed information, that represents the causal structure.

Definition 3.7: For a set of random processes $\underline{\mathbf{X}}$, the *directed information graph* is a directed graph where each node represents a process and there is a directed edge from process $\mathbf{X} \equiv \mathbf{X}_k$ to process $\mathbf{Y} \equiv \mathbf{X}_i$ (for $i, k \in [m]$) iff

$$I(\mathbf{X} \rightarrow \mathbf{Y} \| \underline{\mathbf{X}}_{[m] \setminus \{i,k\}}) > 0.$$

In this model, there is a directed edge from \mathbf{X} to \mathbf{Y} if and only if causal knowledge of \mathbf{X} still influences \mathbf{Y} , even with causal knowledge of all other processes. Since edges are found separately, directed information graphs are unique. Also, directed loops are possible.

Remark 4: Under Assumption 1, since Granger causality and directed information are equivalent in the case of jointly Gaussian processes [51], [52], Granger causality graphs and directed information graphs are the same in that setting.

Minimum generative model graphs and directed information graphs are alternative graphical models to understand the relationship between the past of some processes and the future of others in stochastic dynamical systems. We now investigate in what ways these perspectives are related.

Theorem 3.8: For any joint distribution $P_{\mathbf{X}}$ satisfying Assumption 1, the corresponding minimal generative model graph and directed information graph are equivalent.

The proof is in Appendix B. Markov and Bayesian networks are different graphs, showing different aspects of the structure. That directed information graphs and minimum generative model graphs are the same suggests that they show fundamental causal dependence structures of networks of processes.

IV. STRUCTURAL PROPERTIES OF DIRECTED INFORMATION GRAPHS

The construction methodologies of minimum generative model graphs and directed information graphs are analogous to those of Bayesian and Markov networks, respectively. In this section we describe some analogous properties between the graphical models and investigate relationships between them. We first describe how, analogous to Markov chains, the graphical structure of directed information graphs embodies *causal* Markov chains. We then discuss how the variable dependence structure induced by directed information graphs is the unique dynamic Bayesian network. We also examine causal I-map properties of directed information graphs.

A. Causal Markov chains

In Markov networks, there is an important relationship between a node and its immediate neighbors. Consider a variable Y and its neighbor set \underline{A} , which is called the *Markov boundary*. Any subset $\underline{B} \subseteq \underline{V} \setminus \{Y\}$ of variables containing \underline{A} is called a *Markov blanket*. In Figure 5(b), $\{W\}$ is the Markov boundary for Z . $\{W\}$, $\{W, X\}$, $\{W, Y\}$ and $\{W, Y, X\}$ are the possible Markov blankets.

For each of its neighbors $X \in \underline{A}$, $I(X; Y | \underline{V} \setminus \{X, Y\}) > 0$ holds by construction. This is a pairwise, global test. Furthermore, let \underline{W} any subset of $\underline{V} \setminus \{Y\}$. The Markov chain $\underline{W} - \underline{B} - Y$ also holds and

$$I(\underline{W}; Y) \leq I(\underline{B}; Y)$$

with equality iff $\underline{A} \subseteq \underline{W}$. This follows from the data-processing inequality [61].

There is an analogous phenomenon for minimum generative model graphs. In constructing a minimum generative model (Definition 3.5), for each process \mathbf{X}_i , a unique, minimal set $A(i)$ is found that renders \mathbf{X}_i causally conditionally independent of all other processes. Note that (16) is equivalent to

$$I(\underline{\mathbf{X}}_{[m] \setminus \{i\}} \setminus A(i) \rightarrow \mathbf{X}_i \| \underline{\mathbf{X}}_{A(i)}) = 0. \quad (22)$$

The set $A(i)$ indexes the parents of \mathbf{X}_i in the graph. We will call $\underline{\mathbf{X}}_{A(i)}$ the *causal Markov boundary* of \mathbf{X}_i . The causal Markov boundary contains all of the causal influence pertaining to \mathbf{X}_i , and conditioning on the whole network does not provide any more information. For any subset $B(i) \subseteq [m] \setminus \{i\}$ containing $A(i)$, we call $\underline{\mathbf{X}}_{B(i)}$ a *causal Markov blanket*. Analogous to Markov blankets, causal Markov blankets form causal Markov chains.

Lemma 4.1: Let $P_{\mathbf{X}}$ be a distribution satisfying Assumption 1, \mathbf{X}_i a process and $\underline{\mathbf{X}}_{A(i)}$ its causal Markov blanket. For any subset $B(i)$ with $A(i) \subseteq B(i) \subseteq [m] \setminus \{i\}$ and any subset $W(i) \subseteq [m] \setminus \{i\}$, the causal Markov chain $\underline{\mathbf{X}}_{W(i)} \rightarrow \underline{\mathbf{X}}_{B(i)} \rightarrow \mathbf{X}_i$ holds and

$$I(\underline{\mathbf{X}}_{W(i)} \rightarrow \mathbf{X}_i) \leq I(\underline{\mathbf{X}}_{B(i)} \rightarrow \mathbf{X}_i), \quad (23)$$

with equality iff $A(i) \subseteq W(i)$. The proof is in Appendix C.

By Theorem 3.8, since minimum generative model graphs are equivalent to directed information graphs, for each process \mathbf{X}_i , the set of parent processes are the same, it follows that the method of directed information graphs separately finds each member of the causal Markov boundaries.

B. Variable dependence structure of directed information graphs

In Section III-A, we developed minimum generative models by factorizing the joint distribution over time. By Assumption 1, the next step of the processes became decoupled conditioned on the full past, analogous to the decoupled differential equations in the deterministic case (13). By further removing unnecessary dependencies on other processes, we reduced the factorization (14). That reduction is sufficient for characterizing the unique minimum generative model graphs, which elucidates the temporal dependency structure relating the past of some processes and future of others. We can also examine the dependency structure of the underlying variables, resulting from this factorization and reduction.

Consider the stochastic dynamical system in Example 2 and Figure 6. (12) shows the reduced factorization of the system. The reduction was only over the processes conditioned on. Further reductions are often possible. For instance, if the system was Markov of order one, so the state of the system at time each time $j+1$ only depended on the state at the previous time j , then (12) can be further reduced to

$$\begin{aligned} P_{X^n, Y^n, Z^n}(dx^n, dy^n, dz^n) \\ = \prod_{j=1}^n P_{X_j|X_{j-1}, Y_{j-1}}(dx_j|x_{j-1}, y_{j-1}) \\ \quad \times P_{Y_j|X_{j-1}, Y_{j-1}}(dy_j|x_{j-1}, y_{j-1}) \\ \quad \times P_{Z_j|Y_{j-1}, Z_{j-1}}(dz_j|y_{j-1}, z_{j-1}). \end{aligned} \quad (24)$$

Although this is a significant reduction in variable dependencies, the causal dependence structure between processes is still represented by Figure 6. The variable dependence structure is shown in Figure 7 without any reduction, with removal of unnecessary processes, and with removal of all unnecessary variables. By reducing all unnecessary variables, we obtain a Bayesian network.

Lemma 4.2: For any distribution $P_{\mathbf{X}}$ satisfying Assumption 1, if each term in the factorization (14) for the minimum generative model is fully reduced, the resulting structure of underlying variable dependencies is the unique Bayesian network with causal ordering. The proof is in Appendix D.

Since there is a unique causal ordering for any joint distribution satisfying Assumption 1, once the distribution is factorized over time (13), the unnecessary variables can be removed in any order to obtain the Bayesian network of the underlying variable structure. That is, whole processes need not be removed from causal conditioning at once.

C. Causal I-map properties of directed information graphs

The edges of a Markov network are defined by global conditional dependencies, where the conditioning is on the rest of the network. Additionally, local conditional dependencies, such as conditioning on only a small subset of other nodes, can be identified through simple graphical separation. For instance, for Figure 5(b), while $I(Y; Z|X, W)$ is known from construction, since the only path between Y and Z goes through W , we can further conclude that $I(Y; Z|W) = 0$ by graph separation.

Let \underline{U} , \underline{W} , and \underline{Z} be three disjoint subsets of \underline{V} in an undirected graph. If every path between a node in \underline{U} and a node in \underline{W} contains a node in \underline{Z} , then this is denoted as $\langle \underline{U} | \underline{Z} | \underline{W} \rangle$. For example, in Figure 5(b), $\langle Z | W | Y \rangle$ holds while $\langle X | W | Y \rangle$ does not. An undirected graph is an *independency map* (I-map) if, for all disjoint subsets \underline{U} , \underline{W} , and \underline{Z} ,

$$\langle \underline{U} | \underline{Z} | \underline{W} \rangle \Rightarrow \underline{U} \perp \underline{W} | \underline{Z}.$$

That is, each local or global graphical separation corresponds to a statistical independence relationship. Note that a complete graph, which has no graphical separation, is a trivial I-map. A graph is called a *perfect map* if \Leftarrow also holds.

Markov networks are unique, minimal I-maps [42]. That is, removing any edge renders it not an I-map. Not all systems have a perfect map. For Example 1, $W \perp X$, although they share an edge in the Markov network (Figure 5(b)).

Bayesian networks are also minimal I-maps [42], though not unique as they depend on variable ordering. Also, they have a different graphical separation criterion known as d-separation. \underline{Z} d-separates \underline{U} from \underline{W} , denoted as $\langle \underline{U} | \underline{Z} | \underline{W} \rangle_d$, if for every path (not necessarily directed) between a node in \underline{U} and a node in \underline{W} , there is a node ν such that either

- 1) ν has both edges directed inward and neither ν nor any of ν 's descendants are in \underline{Z} or
- 2) ν does not have both edges directed inward and $\nu \in \underline{Z}$.

For example, in Figure 5(a), $\langle Z | W | Y \rangle_d$ and $\langle X | \emptyset | W \rangle_d$ both hold, though $\langle X | Y | W \rangle_d$ does not, as X and W are conditionally dependent on the common child, Y .

Analogously, local causally conditioned causal independences in a directed information graph can be determined through a graphical separation criterion. Let “cI-map” stand for “causal I-map” and “c-separate” stand for “causally separate.”

Definition 4.3: For a directed information graph for a set of random processes \underline{X} , let \underline{U} , \underline{W} , and \underline{Z} be three disjoint subsets of processes of \underline{X} . \underline{Z} c-separates \underline{U} from \underline{W} , denoted as $\langle \underline{U} || \underline{Z} || \underline{W} \rangle$, if, for every path between a node in \underline{U} and a node in \underline{W} there is a node in $\underline{Z} \cup \underline{W}$ with an outgoing arrow.

C-separation $\langle \underline{X} || \underline{Z} || \underline{Y} \rangle$ represents that \underline{Z} blocks influence from \underline{X} to \underline{Y} . Specifically, for all time j , Y_j is independent of X^{j-1} given $\{Y^{j-1}, Z^{j-1}\}$. In terms of path separation, this criterion is similar to d-separation. However, unlike separation criteria for Markov or Bayesian networks, c-separation is not symmetric,

$$\langle \underline{X} || \underline{Z} || \underline{Y} \rangle \not\Leftarrow \langle \underline{Y} || \underline{Z} || \underline{X} \rangle.$$

Example 4: Consider a system of five processes $\{\mathbf{A}, \dots, \mathbf{E}\}$ with a directed information graph depicted in Figure 8. Examples of c-separation include

- $\langle \mathbf{D} || \mathbf{A} || \mathbf{B} \rangle$ similar to d-separation.
- $\langle \mathbf{D} || \{\mathbf{A}, \mathbf{C}\} || \mathbf{B} \rangle$ unlike d-separation, as \mathbf{C} is a common child.
- $\langle \mathbf{C} || \mathbf{D} || \mathbf{A} \rangle$ unlike d-separation, there is no conditioning on \mathbf{B} .
- $\langle \mathbf{E} || \emptyset || \mathbf{C} \rangle$ unlike d-separation.

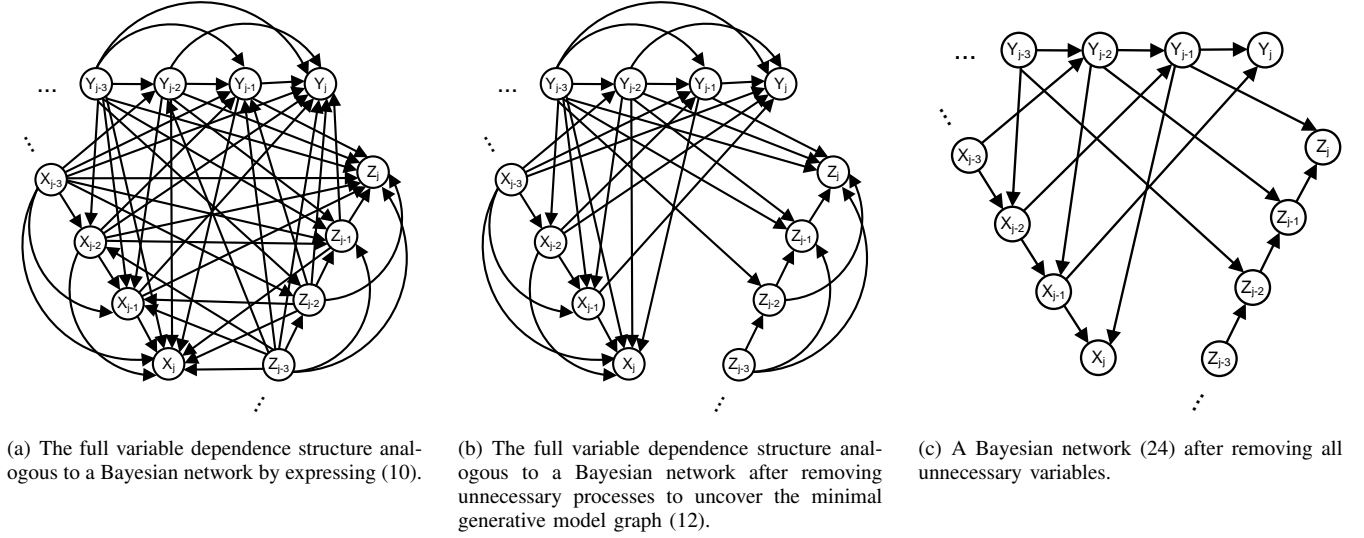


Fig. 7. The variable dependence structure for Example 2 with no reduction (10), removing unnecessary process dependencies (12), and with full reduction (24). Only dependences between variables at time $j - 3$ through j are depicted. Figure (c) is part of the Bayesian network.




Fig. 8. Directed information graph of the system in Example 4.

Remark 5: In addition to asymmetry, another important difference between d-separation and c-separation is dependence on common children. For d-separation, conditioning on common children can render the parents dependent. Bayesian networks depend on the variable indexing. A common child for one index ordering could be a common parent for the reverse index ordering. With directed information graphs, however, causally conditioning on common children (which are not also common ancestors) will not render their parents causally dependence.

We can now define cI-maps for directed graphs.

Definition 4.4: A directed graph for a set of processes \underline{X} is called a *cI-map* if, for all disjoint subsets \underline{U} , \underline{W} , and \underline{Z} of \underline{X} , each c-separation corresponds to a causal Markov chain,

$$\langle \underline{U} \parallel \underline{Z} \parallel \underline{W} \rangle \Rightarrow \underline{U} \rightarrow \underline{Z} \rightarrow \underline{W}.$$

A directed graph is a *minimal* cI-map if removing any of its edges renders it not a valid cI-map.

Theorem 4.5: For any distribution $P_{\underline{X}}$ satisfying Assumption 1, the directed information graph is a minimal cI-map.

For the proof, see Appendix E. Since directed information graphs are minimal cI-maps a natural question is whether they are also perfect maps. That is, in addition to the global causally conditioned independences which are all depicted, are also all local ones depicted? The following example using the

noisy XOR function shows this is not necessarily the case. Causally conditioning on other processes can induce either causal dependence or causal independence.

Example 5: Let W , X , Y , and Z be four processes, with W and X independent Bernoulli($\frac{1}{2}$) processes and

$$\begin{aligned} Y_i &= W_{i-1} \oplus X_{i-1} + \xi_i \\ Z_i &= W_{i-2} \oplus X_{i-2} + \xi'_i \end{aligned}$$

for some i.i.d. Gaussian noises $\{\xi_i, \xi'_i\}_{i=1}^n$. Because of the properties of the XOR function \oplus ,

$$\begin{aligned} I(X \rightarrow Z) &= 0 & \text{but} & & I(X \rightarrow Z \parallel W, Y) &> 0, \\ I(Y \rightarrow Z) &> 0 & \text{but} & & I(Y \rightarrow Z \parallel W, X) &= 0. \end{aligned}$$

We have introduced directed information graphs and minimum generative model graphs as two equivalent procedures for identifying the causal influence structure in a network of processes. To identify the influences, both procedures calculate divergences which use the full joint statistics. We now discuss a procedure for identifying the structure using minimal-dimension statistics when there is some knowledge about the structure.

V. IDENTIFYING THE STRUCTURE

In this section, we discuss methods for identifying the minimal generative model graph - or equivalently, the directed information graph. The methods will take as inputs causally conditioned directed information values. Efficiency will correspond to the dimension of the statistics that will be necessary, such as needing pairwise statistics as compared to the full joint distribution. In particular, we will show that although the definitions of minimum general models and directed information graphs require the full joint statistics, when the number of causal parents of a process Y in the graph is bounded above by K , there is an algorithm to identify Y 's parents using only $(K + 1)$ -wise statistics. This algorithm was

introduced in preliminary work in [62]. We first examine the algorithms for constructing minimum generative model graphs and directed information graphs.

A. General Structures

Identifying the minimum generative model graph by the Definition 3.5 involves determining, for each process \mathbf{X}_i , the minimal cardinality set $A(i)$ that satisfies (16). No search order is prespecified. Since the goal is to find the smallest $A(i)$, one approach is to test increasing sizes of subsets of potential parents. For instance, first the empty set \emptyset is tested, then individual processes, then pairs of processes, etc. This would require calculating an exponential number of causally conditioned directed informations (22).

An alternative method is motivated by causal Markov chains and Lemma 4.1. To find process \mathbf{Y} 's parents, start with the subset of all other processes as a trivial causal Markov blanket and sequentially test each process, shrinking it down into the causal Markov boundary. This method is formally described in Algorithm 1.

Let \mathcal{DI}_{MGM} denote the set of all causally conditioned directed information values from one process to another, causally conditioned on a subset of the the rest:

$$\mathcal{DI}_{MGM} = \left\{ I(\mathbf{X}_k \rightarrow \mathbf{X}_i \parallel \underline{\mathbf{X}}_{B(i)}) : k, i \in [m], B(i) \subseteq [m] \setminus \{i, k\} \right\}.$$

Algorithm 1. MGMconstruct

Input: \mathcal{DI}_{MGM}

1. **For** $i \in [m]$
2. $A(i) \leftarrow [m] \setminus \{i\}$
3. **For** $k \in A(i)$
4. $B(i) \leftarrow A(i) \setminus \{k\}$
5. **If** $I(\mathbf{X}_k \rightarrow \mathbf{X}_i \parallel \underline{\mathbf{X}}_{B(i)}) = 0$
6. $A(i) \leftarrow B(i)$

Lemma 5.1: Let $P_{\underline{\mathbf{X}}}$ be a distribution satisfying Assumption 1. Algorithm 1 recovers the minimal generative model graph.

The proof follows from Lemma 4.1. Algorithm 1 requires the full joint statistics. However, it only uses $\mathcal{O}(m^2)$ tests. Note that the tests used in line 5 are adaptive, using the current $B(i)$. We next consider the algorithm for constructing directed information graphs.

Directed information graphs are identified by testing each edge separately. Testing an edge entails computing all causally conditioned directed informations from one process to another, causally conditioned on all other processes. This is described in Algorithm 2. Let \mathcal{DI}_{DI} denote that set of causally conditioned directed informations:

$$\mathcal{DI}_{DI} = \{ I(\mathbf{X}_k \rightarrow \mathbf{X}_i \parallel \underline{\mathbf{X}}_{[m] \setminus \{i, k\}}) : i, k \in [m] \}.$$

Lemma 5.2: Let $P_{\underline{\mathbf{X}}}$ be a distribution satisfying Assumption 1. Algorithm 2 recovers the directed information graph.

The proof follows from Definition 3.7. Unlike Algorithm 1, Algorithm 2 uses each of the $\mathcal{O}(m^2)$ elements in \mathcal{DI}_{DI} regardless of the graph structure. Line 4 could be executed

Algorithm 2. DIconstruct

Input: \mathcal{DI}_{DI}

1. **For** $i \in [m]$
2. $A(i) \leftarrow \emptyset$
3. **For** $i, j \in [m]$
4. **If** $I(\mathbf{X}_k \rightarrow \mathbf{X}_i \parallel \underline{\mathbf{X}}_{[m] \setminus \{i, k\}}) > 0$
5. $A(i) \leftarrow A(i) \cup \{k\}$




Fig. 9. A graph depicting the task of identifying the parent(s) of \mathbf{Y} the directed information graph in Example 6. It is known that \mathbf{Y} has at most two parents. The structure of $\{\mathbf{X}_1, \dots, \mathbf{X}_5\}$ is not depicted.

in parallel for every possible edge. The number of causally conditioned directed information tests is the same as Algorithm 2, though the tests themselves are different. Note that $\mathcal{DI}_{DI} \subset \mathcal{DI}_{MGM}$.

B. Bounded In-degree

Algorithms 1 and 2 both require inputs computed using the whole joint distribution of the system. Example 5 showed that using the full joint statistics is necessary in general, since if a process has every other process as a parent (for some i , $A(i) = [m] \setminus \{i\}$), the full joint statistics is required to correctly identify this. However, we will show that when there are upper bounds on the in-degrees of the processes, lower-dimension statistics can be used. First consider a simple example.

Example 6: Let $\underline{\mathbf{X}} = \{\mathbf{X}_1, \dots, \mathbf{X}_5, \mathbf{Y}\}$ be a set of six processes with a known joint distribution $P_{\underline{\mathbf{X}}}$ satisfying Assumption 1, and \mathbf{Y} has at most two parents in the minimal generative model structure. Our goal is to identify the parent(s) of \mathbf{Y} . This task is depicted in Figure 9, as any of the other processes could potentially be a parent.

To determine \mathbf{Y} 's parent(s), first we will use Algorithm 1. To test if \mathbf{X}_3 is a parent, calculate $I(\mathbf{X}_3 \rightarrow \mathbf{Y} \parallel \underline{\mathbf{X}}_{\{1,2,4,5\}})$. If this value is 0, then to test next if \mathbf{X}_5 is a parent, calculate $I(\mathbf{X}_5 \rightarrow \mathbf{Y} \parallel \underline{\mathbf{X}}_{\{1,2,4\}})$. Otherwise use $I(\mathbf{X}_5 \rightarrow \mathbf{Y} \parallel \underline{\mathbf{X}}_{\{1,2,3,4\}})$. This process continues until \mathbf{X}_1 , \mathbf{X}_2 , and \mathbf{X}_4 are also tested.

Alternatively, Algorithm 2, the definition of directed information graphs could be used. This allows us to find each parent separately. To test if \mathbf{X}_3 is a parent, calculate $I(\mathbf{X}_3 \rightarrow \mathbf{Y} \parallel \underline{\mathbf{X}}_{\{1,2,4,5\}})$. To test if \mathbf{X}_5 is a parent, calculate $I(\mathbf{X}_5 \rightarrow \mathbf{Y} \parallel \underline{\mathbf{X}}_{\{1,2,3,5\}})$. These tests can be done in parallel.

Both of these methods require the full joint distribution. Given the knowledge that a node has at most K parents in the graph, a natural question is if the full joint distribution is necessary. From Example 5, it is known that at least $(K+1)$ -wise statistics is necessary. We now describe such a method which identifies parents in a distributed manner only using $(K+1)$ -wise statistics.

For each process \mathbf{X}_i , only its causal parents carry the influence, so among all sets of $K(i)$ other processes, all and only those that contain \mathbf{X}_i 's parents will have maximal directed information to \mathbf{X}_i . We can then take the intersection of these sets (causal Markov blankets) to get precisely \mathbf{X}_i 's causal parents (the causal Markov boundary). Algorithm 3 formally describes this method. Let \mathcal{DI}_{BndInd} denote a set of directed information values, such that for each $i \in [m]$, \mathcal{DI}_{BndInd} contains directed information values from each $K(i)$ sized subset of processes to \mathbf{X}_i .

$$\mathcal{DI}_{BndInd} = \left\{ I(\underline{\mathbf{X}}_{B(i)} \rightarrow \mathbf{X}_i) : i \in [m], \right. \\ \left. B(i) \subseteq [m] \setminus \{i\}, |B(i)| = K(i) \right\}.$$

Algorithm 3. StructureRecovery

Input: \mathcal{DI}_{BndInd}

1. **For** $i \in [m]$
2. $A(i) \leftarrow \emptyset$
3. $L \leftarrow K(i)$
4. $\mathcal{I}_L \leftarrow \{\mathcal{I} : \mathcal{I} \subseteq [m] \setminus \{i\}, |\mathcal{I}| = L\}$
5. **For** $\mathcal{I}_l \in \mathcal{I}_L$
6. Compute $I(\underline{\mathbf{X}}_{\mathcal{I}_l} \rightarrow \mathbf{X}_i)$
7. $\mathcal{I}_{Lmax} \leftarrow \arg \max_{\mathcal{I}_j \in \mathcal{I}_L} I(\underline{\mathbf{X}}_{\mathcal{I}_j} \rightarrow \mathbf{X}_i)$
8. $A(i) \leftarrow \bigcap_{\mathcal{I}_l \in \mathcal{I}_{Lmax}} \mathcal{I}_l$

Theorem 5.3: Algorithm 3 recovers the minimal generative model structure for a given $P_{\underline{\mathbf{X}}}$ if for each $i \in [m]$, $K(i) \leq m - 2$.

The proof follows from Lemma 4.1. Note $m - 1$ is a trivial upper bound, as that is the size of the set of all other processes, so there is only one candidate set.

Algorithm 3 finds the structure using only statistics of the dimension of the bound of the in-degree. Thus the bound of the in-degree is both necessary and sufficient for recovering the structure. Algorithm 3 uses all of the elements in \mathcal{DI}_{BndInd} , which are $\sum_{i=1}^m \binom{m-1}{K(i)}$ values. Note that if the upper bounds $\{K(i)\}_{i=1}^m$ do not grow with m , then the algorithm performs $\mathcal{O}(m^{K+1})$ directed information tests, where $K = \max_{i \in [m]} K(i)$. While more tests are used than in Algorithms 1 and 2, since only $(K+1)$ -wise statistics are used, the time to compute or estimate the causally conditioned directed information values for \mathcal{DI}_{BndInd} could be significantly less than that for \mathcal{DI}_{MGM} or \mathcal{DI}_{DI} .

Remark 6: We note that this is not the only algorithm that will recover the structure of trees ($K = 1$ for all nodes) using only pairwise statistics. There are methods [18], [19], analogous to the Chow and Liu algorithm for variables [17], which identify the best tree approximation for causally conditional dependence structures. They compute the directed information between all pairs of processes. However, they find the maximum weight directed spanning tree to determine the best approximating structure, thus coupling identification of parents for different processes. A comparison of the properties of these tree approximation algorithms and Algorithms 1, 2, and 3 (with in-degree bound $K = 1$) is shown in Table 1.

	Distributed Search	Pairwise Statistics
Chow and Liu		X
Alg. 1 and 2	X	
Alg. 3 ($K = 1$)	X	X

Table 1. A comparison of properties of Chow and Liu based algorithms [18], [19] and Algorithms 1, 2, and 3 ($K = 1$). *Distributed search* means that the algorithm finds the parents of a process independently of finding the parents of other processes.

We next simulate a network of stochastic processes and use this procedure to efficiently find the causal parents of each process.

VI. SIMULATION

In this section, we illustrate directed information graphs and the efficient inference method, Algorithm 3, with a simulation. A network of stochastic processes is simulated using a generative model. The directed information graph is then inferred using Algorithm 3 for bounded in-degree. For comparison, the definition of directed information graphs, Algorithm 2, is also used for inference.

A. Setup

1) *Simulation design:* A stochastic network of six coupled point processes $\{\mathbf{A}, \dots, \mathbf{F}\}$ was simulated. We discretized time, using bins of length $\Delta = 1$ ms. The point process model thus becomes


$$\log P(Y_j = y_j | Y^{j-1} = y^{j-1}, X^{j-1} = x^{j-1}) \\ = y_j \log \lambda_j - \lambda_j \Delta, \quad (25)$$

where $y_j \in \{0, 1\}$ and λ_j is $\sigma(Y^{j-1}, X^{j-1})$ -measurable, termed the *conditional intensity function* [63]. We used a generalized linear model with a Poisson link parameter. This class of probability models has a simple analytic form and have been shown to model point processes well, such as the communication between neurons [63]. An example of this distribution for a process \mathbf{Y} at time j causally depending on itself and a process \mathbf{X} for the past L time-steps is characterized by:

$$\log \lambda_j = \alpha_0 + \sum_{l=1}^L \alpha_l y_{j-l} + \beta_l x_{j-l}$$

for constants α_0 and $\{\alpha_l, \beta_l\}_{l=1}^L$. In the simulation, each process at time j explicitly depended on its own past and the past of its parents for six time-steps $\{j-1, \dots, j-6\}$. The parameters $\{\alpha_l, \beta_l\}$ were generated randomly with a normal distribution and fixed for the simulation.

Figure 10(a) depicts the causal dependencies in the generative model used to simulate the network. By the discrete-time construction of the simulation, Assumption 1 holds so the minimum generative model graph and the directed information graph are the same. The network was simulated for ten minutes using one millisecond resolution (6×10^5 time-steps). Although the generative model was known by design, only the simulated data was used to infer the causal influence structure.



(a) The minimum generative model used to simulate the network.

(b) The graph structure inferred using Algorithm 3, assuming an in-degree bound of $K = 3$. Estimation of normalized directed information rate estimates for each K -sized subset were computed. All normalized directed information rate values that were within 5% of that maximum value were considered maximal. Edges $B \rightarrow A$ and $C \rightarrow A$ are missing.

(c) The graph structure inferred using Algorithm 2. An edge was accepted if its normalized, causally conditioned directed information rate estimate was greater than 5%. It incorrectly identified $D \rightarrow B$ as an edge.

Fig. 10. The influence structure of the minimum generative model used to simulate the network and the graphs inferred using Algorithms 2 and 3. They both correctly identified the presence/absence of the 30 potential edges with only one and two mistakes respectively.

2) *Directed information estimation*: There are multiple methods that have been proposed to consistently estimate the directed information, under appropriate assumptions. One is a parametric technique which first does model fitting and then takes the empirical average of log likelihoods [34], [35], [37] and can be applicable to point process modalities. The other is a universal estimation technique using the context weighting tree method [41] that is applicable to discrete-time processes with finite alphabets. We employed the parametric technique in [34]. We now sketch the procedure.

The parametric technique estimates the causally conditioned directed information rate [13]

$$\bar{I}(\underline{\mathbf{X}}_{\mathcal{I}_1} \rightarrow \mathbf{Y} \parallel \underline{\mathbf{X}}_{\mathcal{I}_2}) \triangleq \lim_{n \rightarrow \infty} \frac{1}{n} I(\underline{\mathbf{X}}_{\mathcal{I}_1} \rightarrow \mathbf{Y} \parallel \underline{\mathbf{X}}_{\mathcal{I}_2})$$

for given subsets $\mathcal{I}_1, \mathcal{I}_2 \subseteq [m] \setminus \{i\}$. It does so by first estimating two causally conditioned entropy rates

$$\bar{H}(\mathbf{Y} \parallel \underline{\mathbf{X}}_{\mathcal{I}_2}) \triangleq \lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}_{P_{\mathbf{Y}, \underline{\mathbf{X}}_{\mathcal{I}_2}}} [-\log P_{\mathbf{Y} \parallel \underline{\mathbf{X}}_{\mathcal{I}_2}}(\mathbf{Y} \parallel \underline{\mathbf{X}}_{\mathcal{I}_2})]$$

and $\bar{H}(\mathbf{Y} \parallel \underline{\mathbf{X}}_{\mathcal{I}_2 \cup \mathcal{I}_1})$ defined likewise.

Then the causally conditioned directed information rate is their difference

$$\bar{I}(\underline{\mathbf{X}}_{\mathcal{I}_1} \rightarrow \mathbf{Y} \parallel \underline{\mathbf{X}}_{\mathcal{I}_2}) = \bar{H}(\mathbf{Y} \parallel \underline{\mathbf{X}}_{\mathcal{I}_2}) - \bar{H}(\mathbf{Y} \parallel \underline{\mathbf{X}}_{\mathcal{I}_2 \cup \mathcal{I}_1}).$$

The technique assumes $P_{\mathbf{Y}, \underline{\mathbf{X}}_{\mathcal{I}_2 \cup \mathcal{I}_1}}$

- is jointly stationary and ergodic;
- is Markov of some finite order J^* , with J^* unknown but bounded;
- belongs to a known parametric class of distributions.

These assumptions imply $P_{Y_j | Y^{j-1}, \underline{\mathbf{X}}_{\mathcal{I}_2 \cup \mathcal{I}_1}^{j-1}}$ is characterized by a constant parameter vector $\theta \in \mathbb{R}^{J^*}$ for all time j .

The procedure estimates causally conditioned entropies as follows [34]:

- 1) do maximum likelihood model fits to estimate $\bar{\theta}(J)$ for various model orders J ;
- 2) use a model order selection criterion [64] to determine the best-fitting model order \hat{J}^* ;

- 3) compute the empirical average of the log likelihood of $P_{Y_j | Y^{j-1}, \underline{\mathbf{X}}_{\mathcal{I}_2 \cup \mathcal{I}_1}^{j-1}}$ characterized by $\hat{\theta}(\hat{J}^*)$. This average is the estimate $\hat{H}(\mathbf{Y} \parallel \underline{\mathbf{X}}_{\mathcal{I}_2 \cup \mathcal{I}_1})$ for the causally conditional entropy $\bar{H}(\mathbf{Y} \parallel \underline{\mathbf{X}}_{\mathcal{I}_2 \cup \mathcal{I}_1})$.

Instead of using the causally conditioned directed information rate estimates directly, normalized values were used

$$\hat{I}(\underline{\mathbf{X}}_{\mathcal{I}_1} \rightarrow \mathbf{Y} \parallel \underline{\mathbf{X}}_{\mathcal{I}_2}) \triangleq \frac{\hat{H}(\mathbf{Y} \parallel \underline{\mathbf{X}}_{\mathcal{I}_2}) - \hat{H}(\mathbf{Y} \parallel \underline{\mathbf{X}}_{\mathcal{I}_2 \cup \mathcal{I}_1})}{\hat{H}(\mathbf{Y} \parallel \underline{\mathbf{X}}_{\mathcal{I}_2})}.$$

In implementing this procedure to analyze the simulated data, it was assumed that $P_{Y_j | Y^{j-1}, \underline{\mathbf{X}}_{\mathcal{I}_2 \cup \mathcal{I}_1}^{j-1}}$ belonged to the family of generalized linear models with Poisson link function (25). The model fits were computed using Matlab function `glmfit(.)` for a fixed model order $J = 10$.

3) *Graph structure identification*: First, the structure was inferred using Algorithm 2, which follows from the definition of directed information graphs. Each of the 30 possible edges was tested separately. To test the edge $A \rightarrow B$, for instance, $\hat{I}(A \rightarrow B \parallel \{C, D, E\})$ was computed. Other edges were tested likewise. An edge was accepted if its normalized, causally conditioned directed information rate estimate was greater than 5%.

Additionally, Algorithm 3 was used under an assumed upper bound of $K = 3$. This value is a strict upper-bound for the network, though only one process had three parents. Thus, for each process, $\binom{5}{3} = 10$ sets of candidate parents were considered. To find the parents of process A , for example, the normalized directed information rate estimates for each K -sized subset of $\{B, C, D, E\}$ to A were computed. The maximum normalized directed information rate was identified. Then all other normalized directed information rate values that were within 5% of that maximum value were considered maximal. The intersection of the corresponding maximal subsets was taken to identify the parents.

B. Results

The structure of the simulated network inferred using Algorithm 3 is shown in Figure 10(b). This is the same structure as the generative model, except that it is missing $B \rightarrow A$ and $C \rightarrow A$. For this simulation, most of the parent sets had distinguishably strong influence. For instance, for inferring the parents of process C , the sets $\{A, D, F\}$, $\{B, D, F\}$, and $\{E, D, F\}$ all had normalized directed information rate values of 62%, with the next highest value of 51% from set $\{A, D, E\}$. However, for inferring the parents of A , $\{B, C, E\}$ had a estimated value of 57%, while $\{B, D, E\}$ and $\{C, E, F\}$ had values of 52%.

The structure inferred using Algorithm 2, the definition of directed information graphs, is shown in Figure 10(c). It is almost correct. However, this algorithm additionally detected the edge $D \rightarrow B$. The normalized directed information of that edge was close to threshold, $\hat{I}(D \rightarrow B | A, C, E, F) = 6\%$. All of the other inferred edges corresponded to higher normalized values, from 11% to 51%. While this method computed fewer directed information estimates (30 as opposed to 60), it required estimating full joint statistics.

VII. CONCLUSION AND FUTURE DIRECTIONS

Methods that characterize the causal dependence structure of a network of dynamically interacting processes could significantly bolster research in a number of diverse disciplines, including social sciences, economics, biology, and physics. We propose two graphical models that represent the causal dependence structure of such networks. One is motivated by generative models and the other by Granger causality. Despite their different perspectives, the structures they depict are equivalent. The presence or absence of an edge is determined by calculating an information theoretic divergence known as directed information.

There are a number of directions for future research to facilitate applying these methods to real-world data. One direction is improving estimation techniques of causally conditioned directed information. As discussed in this paper, there are already several estimation techniques [34], [35], [37], [41]. Computational feasibility of current methods is an aspect that will need to be further explored, especially for data-rich applications. For instance, social networks could involve millions of users over the course of several years. It will be important to improve current estimation techniques and develop new ones to efficiently identify the causal dependence structure even of such large datasets.

Another direction of future research involves extending these graphical models to time-varying network structures. The graphical models proposed here model dynamically interacting processes, but assume the causal dependence structure itself is time-invariant. For a number of real-world networks, the network influence structure varies with time, such as the internet, the brain, and social networks like Twitter. Incorporating these future research extensions could provide for a practical framework to analyze a variety of real-world networks of causally interacting processes.




Fig. 11. Sets of processes for Lemma A.1.

APPENDIX A PROOF OF LEMMA 3.4

For the proof that minimum generative models are unique, namely that the $A(i)$'s can be selected independently and there is a unique one of minimum size, we use the following lemma. It is a variation of the *intersection property* which is used in the proof of uniqueness for Bayesian networks for a given ordering ([42] pp. 119). Essentially, the lemma says that no two subsets of processes can influence a process Y in exactly the same way, unless all the influence comes from their intersection.

Lemma A.1: Let $P_{\mathbf{X}}$ denote a joint distribution satisfying Assumption 1. Let $\mathbf{Y} \equiv \mathbf{X}_i$, $\mathbf{U} \equiv \mathbf{X}_{\mathcal{I}_{\mathbf{U}}}$, and $\mathbf{W} \equiv \mathbf{X}_{\mathcal{I}_{\mathbf{W}}}$ such that $\mathcal{I}_{\mathbf{U}}, \mathcal{I}_{\mathbf{W}} \subseteq [m] \setminus \{i\}$. Denote the intersection as $\mathbf{V} = \mathbf{U} \cap \mathbf{W}$. (See Figure 11.)

$$\text{If } D(P_{\mathbf{Y}|\mathbf{U}\mathbf{W}} \| P_{\mathbf{Y}|\mathbf{U}} | P_{\mathbf{U}\mathbf{W}}) = 0 \quad \text{and} \quad (26)$$

$$D(P_{\mathbf{Y}|\mathbf{U}\mathbf{W}} \| P_{\mathbf{Y}|\mathbf{W}} | P_{\mathbf{U}\mathbf{W}}) = 0, \quad (27)$$

$$\text{then } D(P_{\mathbf{Y}|\mathbf{U}\mathbf{W}} \| P_{\mathbf{Y}|\mathbf{V}} | P_{\mathbf{U}\mathbf{W}}) = 0. \quad (28)$$

Proof: Since $P_{\mathbf{X}}$ is positive, the divergences (26), (27), and (28) are equivalent to conditional independence statements of the underlying random variables. For instance, using marginalization and non-negativity of KL divergence, (26) is equivalent to

$$Y_j \perp\!\!\!\perp \mathbf{W}^{j-1} \setminus \mathbf{V}^{j-1} \mid \mathbf{U}^{j-1}, Y^{j-1}$$

for all time $1 \leq j \leq n$. The proof then directly follows from the intersection property ([42] pp. 84). ■

We can now prove Lemma 3.4.

Proof: Suppose not. Let A and B be distinct two minimal generative models for $P_{\mathbf{X}}$. Let $\mathbf{Y} \equiv \mathbf{X}_i$ for some $i \in [m]$ be any process for which $A(i) \neq B(i)$. By definition of minimal generative models, properties of the logarithm, and non-negativity of KL-divergence,

$$D(P_{\mathbf{Y}|\mathbf{X}_{[m] \setminus \{i\}}} \| P_{\mathbf{Y}|\mathbf{X}_{A(i)}} | P_{\mathbf{X}_{[m] \setminus \{i\}}}) = 0 \quad \text{and}$$

$$D(P_{\mathbf{Y}|\mathbf{X}_{[m] \setminus \{i\}}} \| P_{\mathbf{Y}|\mathbf{X}_{B(i)}} | P_{\mathbf{X}_{[m] \setminus \{i\}}}) = 0.$$

Thus, by Lemma A.1,

$$D(P_{\mathbf{Y}|\mathbf{X}_{[m] \setminus \{i\}}} \| P_{\mathbf{Y}|\mathbf{X}_{A(i) \cap B(i)}} | P_{\mathbf{X}_{[m] \setminus \{i\}}}) = 0.$$

This is a contradiction, as $|A(i) \cap B(i)| < |A(i)|$ but $|A(i)|$ is minimal by definition. ■

APPENDIX B PROOF OF THEOREM 3.8

Proof: Let $\{A(i)\}_{i=1}^m$ denote the parent sets in the minimal generative model, and let $\mathbf{X} \equiv \mathbf{X}_k$ and $\mathbf{Y} \equiv \mathbf{X}_i$ be processes for some $i, k \in [m]$. We consider two cases. First suppose that there is no edge from \mathbf{X} to \mathbf{Y} in the minimal generative model graph. By definition, $k \notin A(i)$. Thus,

$$D(P_{\mathbf{Y} \parallel \underline{\mathbf{X}}_{[m] \setminus \{i\}}} \parallel P_{\mathbf{Y} \parallel \underline{\mathbf{X}}_{[m] \setminus \{i, k\}}} | P_{\underline{\mathbf{X}}_{[m] \setminus \{i\}}}) = 0, \quad (29)$$

and consequently

$$I(\mathbf{X} \rightarrow \mathbf{Y} \parallel \underline{\mathbf{X}}_{[m] \setminus \{i, k\}}) = 0. \quad (30)$$

This means that there is no edge from \mathbf{X} to \mathbf{Y} in the directed information graph either.

Now suppose that there is no edge from \mathbf{X} to \mathbf{Y} in the directed information graph. Equation (30) holds which implies (29). However, suppose there is a directed edge from \mathbf{X} to \mathbf{Y} in the minimal generative model graph. Thus $k \in A(i)$, and $|A(i)|$ is minimal. With (29) this implies by Lemma A.1

$$D(P_{\mathbf{Y} \parallel \underline{\mathbf{X}}_{[m] \setminus \{i\}}} \parallel P_{\mathbf{Y} \parallel \underline{\mathbf{X}}_{A(i) \setminus \{k\}}} | P_{\underline{\mathbf{X}}_{[m] \setminus \{i\}}}) = 0.$$

This is a contradiction because $|A(i)|$ is minimal. Therefore there is no edge in the minimal generative model graph either. ■

APPENDIX C PROOF OF LEMMA 4.1

The proof is based on applying the chain rule different ways, non-negativity of directed information, and Lemma 3.4.

Proof: First apply the chain rule different ways:

$$\begin{aligned} I(\mathbf{X}_{W(i) \cup B(i)} \rightarrow \mathbf{X}_i) &= I(\mathbf{X}_{W(i)} \rightarrow \mathbf{X}_i) + I(\mathbf{X}_{B(i)} \rightarrow \mathbf{X}_i \parallel \mathbf{X}_{W(i)}) \\ &= I(\mathbf{X}_{B(i)} \rightarrow \mathbf{X}_i) + I(\mathbf{X}_{W(i)} \rightarrow \mathbf{X}_i \parallel \mathbf{X}_{B(i)}) \quad (31) \\ &= I(\mathbf{X}_{A(i)} \rightarrow \mathbf{X}_i) + I(\mathbf{X}_{W(i) \cup B(i)} \rightarrow \mathbf{X}_i \parallel \mathbf{X}_{A(i)}). \quad (32) \end{aligned}$$

Note by definition of $A(i)$ in the minimum generative models,

$$I(\mathbf{X}_{W(i) \cup B(i)} \rightarrow \mathbf{X}_i \parallel \mathbf{X}_{A(i)}) = 0.$$

Next apply the chain rule to $I(\mathbf{X}_{B(i)} \rightarrow \mathbf{X}_i)$ in (31) and combine with (32).

$$\begin{aligned} I(\mathbf{X}_{A(i)} \rightarrow \mathbf{X}_i) &= I(\mathbf{X}_{A(i)} \rightarrow \mathbf{X}_i) \\ &\quad + I(\mathbf{X}_{B(i)} \rightarrow \mathbf{X}_i \parallel \mathbf{X}_{A(i)}) + I(\mathbf{X}_{W(i)} \rightarrow \mathbf{X}_i \parallel \mathbf{X}_{B(i)}). \end{aligned}$$

Consequently, since directed information is a KL divergence and thus nonnegative, $I(\mathbf{X}_{W(i)} \rightarrow \mathbf{X}_i \parallel \mathbf{X}_{B(i)}) = 0$ and

$$I(\mathbf{X}_{W(i)} \rightarrow \mathbf{X}_i) \leq I(\mathbf{X}_{B(i)} \rightarrow \mathbf{X}_i).$$

This is the inequality (23) in Lemma 4.1. Equality occurs when $I(\mathbf{X}_{B(i)} \rightarrow \mathbf{X}_i \parallel \mathbf{X}_{W(i)}) = 0$. This corresponds to the divergence

$$D(P_{\mathbf{X}_i \parallel \mathbf{X}_{W(i) \cup B(i)}} \parallel P_{\mathbf{X}_i \parallel \mathbf{X}_{W(i)}} | P_{\mathbf{X}_{W(i) \cup B(i)}}) = 0.$$

Since we also have that

$$D(P_{\mathbf{X}_i \parallel \mathbf{X}_{W(i) \cup B(i)}} \parallel P_{\mathbf{X}_i \parallel \mathbf{X}_{A(i)}} | P_{\mathbf{X}_{W(i) \cup B(i)}}) = 0,$$

by Lemma 3.4, $A(i) \subseteq W(i)$, else $A(i)$ would not be minimal which is a contradiction. ■

APPENDIX D PROOF OF LEMMA 4.2

First consider the procedure for constructing a Bayesian network, outlined in Section II-B. For each variable X_i , let $\underline{A}_i \subseteq \{X_1, X_2, \dots, X_{i-1}\}$ denote a Markov boundary of X_i with respect to the set of preceding variables $\{X_1, X_2, \dots, X_{i-1}\}$. The directed acyclic graph (DAG) formed by setting each member of \underline{A}_i as a parent of X_i is called a *boundary DAG* with respect to the index order. By [65], boundary DAGs are Bayesian networks (minimal I-maps under d-separation). This allows for a relatively simple procedure (factorization and reduction) to identify a minimal I-map, regardless of the initial ordering. We now prove Lemma 4.2.

Proof: The minimum generative model specifies an ordering to apply the chain rule. The ordering is first over time, and, with Assumption 1, the variables at the same instance are conditionally independent given the past. Consequently, any causal ordering results in the same factorization.

Next, some variable dependencies are removed. Specifically, for each process \mathbf{X}_i , processes not in its causal Markov boundary are unconditioned on for each conditional term $P_{X_{i,j} | \underline{\mathbf{X}}_{[m] \setminus \{i\}}^{j-1}}$. There is still conditioning on the full past of all processes in \mathbf{X}_i 's causal Markov boundary. By further removing all unnecessary variable dependencies for each conditional term, which can be done uniquely since the joint distribution is positive, a boundary DAG, and thus a Bayesian network, for the variable dependence structure is obtained. ■

APPENDIX E PROOF OF THEOREM 4.5

The proof is based on d-separation in the Bayesian network of variable dependencies underlying the directed information graph. We will use the following corollary to simplify the argument.

Corollary E.1: Let $P_{\underline{\mathbf{X}}}$ be any distribution satisfying Assumption 1. Consider its directed information graph and the corresponding induced Bayesian network. Any path between processes in the Bayesian network corresponds to the same path (possibly reusing edges) in the directed information graph.

The proof follows from Lemma 4.2. We now prove Theorem 4.5.

Proof: We first show that $\underline{\mathbf{Z}} \cup \underline{\mathbf{W}}$ d-separates $\underline{\mathbf{U}}$ from the parents of $\underline{\mathbf{W}}$ not in $\underline{\mathbf{Z}}$. We then show this fact implies the result.

Let $\underline{\mathbf{T}}$ denote the set of nodes that have a child in $\underline{\mathbf{W}}$ and are not in $\underline{\mathbf{Z}}$. Any path p from a node in $\underline{\mathbf{U}}$ to a node in $\underline{\mathbf{T}}$ can be extended by one edge into a path from $\underline{\mathbf{U}}$ to $\underline{\mathbf{W}}$. Thus, by definition of c-separation, there must be a node in $\underline{\mathbf{Z}} \cup \underline{\mathbf{W}}$ that is on path p with an outgoing edge.

Thus, we have that in the directed information graph, $\underline{\mathbf{Z}} \cup \underline{\mathbf{W}}$ d-separates $\underline{\mathbf{U}}$ from $\underline{\mathbf{T}}$. By Corollary E.1, this corresponds to d-separation in the Bayesian network. For all time j ,

$$\underline{\mathbf{U}}^{j-1} \perp\!\!\!\perp \underline{\mathbf{T}}^{j-1} \mid \underline{\mathbf{Z}}^{j-1}, \underline{\mathbf{W}}^{j-1}. \quad (33)$$

Since $\underline{\mathbf{Z}} \cup \underline{\mathbf{T}}$ forms a causal Markov blanket for $\underline{\mathbf{W}}$, we have that

$$\underline{\mathbf{U}}^{j-1} \perp\!\!\!\perp \underline{\mathbf{W}}_j \mid \underline{\mathbf{Z}}^{j-1}, \underline{\mathbf{W}}^{j-1}, \underline{\mathbf{T}}^{j-1}. \quad (34)$$

By the contraction property ([42] pg. 84), equations (33) and (34) imply

$$\underline{\mathbf{U}}^{j-1} \perp\!\!\!\perp \{\underline{\mathbf{W}}_j, \underline{\mathbf{T}}^{j-1}\} \mid \underline{\mathbf{Z}}^{j-1}, \underline{\mathbf{W}}^{j-1}.$$

By decomposition,

$$\underline{\mathbf{U}}^{j-1} \perp\!\!\!\perp \underline{\mathbf{W}}_j \mid \underline{\mathbf{Z}}^{j-1}, \underline{\mathbf{W}}^{j-1},$$

or, equivalently, the following causal Markov chain holds:

$$\underline{\mathbf{U}} \rightarrow \underline{\mathbf{Z}} \rightarrow \underline{\mathbf{W}}.$$

This proves that directed information graphs are ci-maps. We next consider minimality.

Consider any two processes \mathbf{X}_i and \mathbf{X}_k in $\underline{\mathbf{X}}$ with an edge $\mathbf{X}_k \rightarrow \mathbf{X}_i$. Removing this edge means $\mathbf{X}_k \parallel \underline{\mathbf{X}}_{[m] \setminus \{i,k\}} \parallel \mathbf{X}_i >$ which implies that $\mathbf{X}_k \rightarrow \underline{\mathbf{X}}_{[m] \setminus \{i,k\}} \rightarrow \mathbf{X}_i$ holds. By construction of the directed information graph, this causal independence statement is incorrect, leading to a contradiction. Thus, directed information graphs are minimal ci-maps. ■

REFERENCES

- [1] E. B. Sudderth, "Graphical models for visual object recognition and tracking," Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, MA, USA, 2006.
- [2] H. Loeliger, "An introduction to factor graphs," *Signal Processing Magazine, IEEE*, vol. 21, no. 1, pp. 28–41, 2004.
- [3] N. Friedman, "Inferring cellular networks using probabilistic graphical models," *Science*, vol. 303, no. 5659, pp. 799–805, 2004.
- [4] L. Andersen, J. Krebs, and J. Andersen, "Steno: An expert system for medical diagnosis based on graphical models and model search," *Journal of Applied Statistics*, vol. 18, no. 1, pp. 139–153, 1991.
- [5] D. Koller and N. Friedman, *Probabilistic graphical models: principles and techniques*. The MIT Press, 2009.
- [6] F. Briggs and W. Usrey, "Corticogeniculate feedback and visual processing in the primate," *The Journal of physiology*, vol. 589, no. 1, pp. 33–40, 2011.
- [7] W. Hesse, E. Möller, M. Arnold, and B. Schack, "The use of time-variant EEG Granger causality for inspecting directed interdependencies of neural assemblies," *Journal of Neuroscience Methods*, vol. 124, no. 1, pp. 27–44, 2003.
- [8] A. Brovelli, M. Ding, A. Ledberg, Y. Chen, R. Nakamura, and S. Bressler, "Beta oscillations in a large-scale sensorimotor cortical network: directional influences revealed by Granger causality," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 101, no. 26, p. 9849, 2004.
- [9] W. Joerding, "Economic growth and defense spending: Granger causality," *Journal of Development Economics*, vol. 21, no. 1, pp. 35–40, 1986.
- [10] C. Hiemstra and J. Jones, "Testing for linear and nonlinear Granger causality in the stock price-volume relation," *Journal of Finance*, pp. 1639–1664, 1994.
- [11] M. Trusov, R. Bucklin, and K. Pauwels, "Effects of word-of-mouth versus traditional marketing: Findings from an internet social networking site," *Journal of Marketing*, vol. 73, no. 5, pp. 90–102, 2009.
- [12] J. Groshek, "Media, instability, and democracy: examining the Granger-causal relationships of 122 countries from 1946 to 2003," *Journal of Communication*, vol. 61, no. 6, pp. 1161–1182, 2011.
- [13] G. Kramer, "Directed information for channels with feedback," Ph.D. dissertation, University of Manitoba, Canada, 1998.
- [14] S. Strogatz, "Exploring complex networks," *Nature*, vol. 410, no. 6825, pp. 268–276, 2001.
- [15] J. Banavar, F. Colaiori, A. Flammini, A. Maritan, and A. Rinaldo, "Topology of the fittest transportation network," *Physical Review Letters*, vol. 84, no. 20, pp. 4745–4748, 2000.
- [16] E. Bullitt, K. Muller, I. Jung, W. Lin, and S. Aylward, "Analyzing attributes of vessel populations," *Medical image analysis*, vol. 9, no. 1, pp. 39–49, 2005.
- [17] C. Chow and C. Liu, "Approximating discrete probability distributions with dependence trees," *IEEE transactions on Information Theory*, vol. 14, no. 3, pp. 462–467, 1968.
- [18] C. Quinn, T. Coleman, and N. Kiyavash, "Causal dependence tree approximations of joint distributions for multiple random processes," *Information Theory, IEEE Transactions on*, 2011, submitted, Arxiv preprint arXiv:1101.5108.
- [19] D. Materassi and G. Innocenti, "Topological identification in networks of dynamical systems," *Automatic Control, IEEE Transactions on*, vol. 55, no. 8, pp. 1860–1871, 2010.
- [20] H. Marko, "The bidirectional communication theory—a generalization of information theory," *Communications, IEEE Transactions on*, vol. 21, no. 12, pp. 1345–1351, Dec 1973.
- [21] J. Rissanen and M. Wax, "Measures of mutual and causal dependence between two time series (Corresp.)," *IEEE Transactions on Information Theory*, vol. 33, no. 4, pp. 598–601, 1987.
- [22] C. Granger, "Investigating causal relations by econometric models and cross-spectral methods," *Econometrica*, vol. 37, no. 3, pp. 424–438, 1969.
- [23] N. Wiener, "The theory of prediction," in *Modern Mathematics for Engineers*, E. F. Beckenback, Ed. New York: McGraw-Hill, 1956.
- [24] J. Massey, "Causality, feedback and directed information," in *Proc. Int. Symp. Information Theory Application (ISITA-90)*, 1990, pp. 303–305.
- [25] S. Tatikonda and S. Mitter, "The Capacity of Channels With Feedback," *IEEE Transactions on Information Theory*, vol. 55, no. 1, pp. 323–349, 2009.
- [26] H. Permuter, T. Weissman, and A. Goldsmith, "Finite state channels with time-invariant deterministic feedback," *IEEE Transactions on Information Theory*, vol. 55, no. 2, pp. 644–662, 2009.
- [27] J. Massey and P. Massey, "Conservation of mutual and directed information," in *IEEE International Symposium on Information Theory, 2005*, 2005, pp. 157–158.
- [28] H. Permuter, Y. Kim, and T. Weissman, "On directed information and gambling," in *IEEE International Symposium on Information Theory, 2008. ISIT 2008*, 2008, pp. 1403–1407.
- [29] —, "Interpretations of directed information in portfolio theory, data compression, and hypothesis testing," *Information Theory, IEEE Transactions on*, vol. 57, no. 6, pp. 3248–3259, 2011.
- [30] N. Elia, "When Bode meets Shannon: control-oriented feedback communication schemes," *Automatic Control, IEEE Transactions on*, vol. 49, no. 9, pp. 1477 – 1488, Sept. 2004.
- [31] N. Martins and M. Dahleh, "Feedback control in the presence of noisy channels: "Bode-like" fundamental limitations of performance," *Automatic Control, IEEE Transactions on*, vol. 53, no. 7, pp. 1604 – 1615, Aug. 2008.
- [32] S. Tatikonda, "Control under communication constraints," Ph.D. dissertation, Massachusetts Institute of Technology, 2000.
- [33] R. Venkataramanan and S. Pradhan, "Source coding with feed-forward: rate-distortion theorems and error exponents for a general source," *IEEE Transactions on Information Theory*, vol. 53, no. 6, pp. 2154–2179, 2007.
- [34] C. Quinn, T. Coleman, N. Kiyavash, and N. Hatsopoulos, "Estimating the directed information to infer causal relationships in ensemble neural spike train recordings," *Journal of computational neuroscience*, vol. 30, no. 1, pp. 17–44, 2011.
- [35] S. Kim, D. Putrino, S. Ghosh, and E. N. Brown, "A Granger causality measure for point process models of ensemble neural spiking activity," *PLoS Comput Biol*, vol. 7, no. 3, 03 2011.
- [36] Y. Liu and S. Aviyente, "Information theoretic approach to quantify causal neural interactions from EEG," in *Signals, Systems and Computers (ASILOMAR), 2010 Conference Record of the Forty Fourth Asilomar Conference on*, nov. 2010, pp. 1380 –1384.
- [37] K. So, A. Koralek, K. Ganguly, M. Gastpar, and J. Carmena, "Assessing functional connectivity of neural ensembles using directed information," *Journal of Neural Engineering*, vol. 9, p. 026004, 2012.
- [38] A. Rao, A. Hero, and J. Engel, "Motif discovery in tissue-specific regulatory sequences using directed information," *EURASIP Journal on Bioinformatics and Systems Biology*, vol. 2007, p. 3, 2007.
- [39] A. Rao, A. Hero, D. States, and J. Engel, "Using directed information to build biologically relevant influence networks," *Journal of bioinformatics and computational biology*, vol. 6, no. 3, pp. 493–520, 2008.
- [40] X. Chen, A. Hero, and S. Savarese, "Multimodal video indexing and retrieval using directed information," *Multimedia, IEEE Transactions on*, vol. 14, no. 1, pp. 3 –16, Feb. 2012.

- [41] J. Jiao, H. Permuter, L. Zhao, Y. Kim, and T. Weissman, "Universal estimation of directed information," *Arxiv preprint arXiv:1201.2334*, 2012.
- [42] J. Pearl, *Probabilistic reasoning in intelligent systems: networks of plausible inference*, 2nd ed. Morgan Kaufmann, 1988.
- [43] K. Murphy, "Dynamic Bayesian networks: representation, inference and learning," Ph.D. dissertation, University of California, 2002.
- [44] J. Pearl, *Causality: models, reasoning and inference*, 2nd ed. Cambridge Univ Press, 2009.
- [45] N. Ay and D. Polani, "Information flows in causal networks," *Advances in Complex Systems*, vol. 11, no. 1, pp. 17–42, 2008.
- [46] M. Raginsky, "Directed information and Pearl's causal calculus," in *Communication, Control, and Computing (Allerton), 2011 49th Annual Allerton Conference on*. IEEE, 2011, pp. 958–965.
- [47] R. Dahlhaus and M. Eichler, "Causality and graphical models in time series analysis," *Oxford Statistical Science Series*, pp. 115–137, 2003.
- [48] R. Dahlhaus, "Graphical interaction models for multivariate time series 1," *Metrika*, vol. 51, no. 2, pp. 157–172, 2000.
- [49] M. Eichler, "Granger causality and path diagrams for multivariate time series," *Journal of Econometrics*, vol. 137, no. 2, pp. 334–353, 2007.
- [50] P. Amblard and O. Michel, "On directed information theory and Granger causality graphs," *Journal of computational neuroscience*, vol. 30, no. 1, pp. 7–16, 2011.
- [51] —, "Relating Granger causality to directed information theory for networks of stochastic processes," *Arxiv preprint arXiv:0911.2873*, 2009.
- [52] L. Barnett, A. Barrett, and A. Seth, "Granger causality and transfer entropy are equivalent for Gaussian variables," *Physical review letters*, vol. 103, no. 23, p. 238701, 2009.
- [53] T. Schreiber, "Measuring information transfer," *Physical review letters*, vol. 85, no. 2, pp. 461–464, 2000.
- [54] P. Spirtes and C. Glymour, "An algorithm for fast recovery of sparse causal graphs," *Social Science Computer Review*, vol. 9, no. 1, p. 62, 1991.
- [55] F. Bromberg, D. Margaritis, and V. Honavar, "Efficient Markov network structure discovery using independence tests," *Journal of Artificial Intelligence Research*, vol. 35, no. 1, pp. 449–484, 2009.
- [56] A. Bolstad, B. Van Veen, and R. Nowak, "Causal network inference via group sparse regularization," *Signal Processing, IEEE Transactions on*, vol. 59, no. 6, pp. 2628–2641, 2011.
- [57] C. Quinn, N. Kiyavash, and T. Coleman, "Equivalence between minimal generative model graphs and directed information graphs," in *Information Theory Proceedings (ISIT), IEEE International Symposium on*, 2011.
- [58] B. Øksendal, *Stochastic differential equations: an introduction with applications*. Springer Verlag, 2003.
- [59] N. Cesa-Bianchi and G. Lugosi, *Prediction, learning, and games*. Cambridge University Press, 2006.
- [60] C. Quinn, T. Coleman, and N. Kiyavash, "A generalized prediction framework for Granger causality," in *INFOCOM*. IEEE, 2011, pp. 906–911, invited Paper for Special Session on Network Science for Communication Networks.
- [61] T. Cover and J. Thomas, *Elements of information theory*. Wiley-Interscience, 2006.
- [62] C. Quinn, N. Kiyavash, and T. Coleman, "A minimal approach to causal inference on topologies with bounded indegree," in *IEEE Conference on Decision and Control*, 2011.
- [63] W. Truccolo, U. Eden, M. Fellows, J. Donoghue, and E. Brown, "A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects," *Journal of Neurophysiology*, vol. 93, no. 2, p. 1074, 2005.
- [64] A. Barron, J. Rissanen, and B. Yu, "The minimum description length principle in coding and modeling," *Information Theory, IEEE Transactions on*, vol. 44, no. 6, pp. 2743–2760, 1998.
- [65] T. Verma and J. Pearl, "Causal networks: Semantics and expressiveness," in *Proceedings of the Proceedings of the Fourth Conference Annual Conference on Uncertainty in Artificial Intelligence (UAI-88)*. New York, NY: Elsevier Science, 1988, pp. 352–359.